Exhibit 19

Filed Under Seal

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

SAMSUNG ELECTRONICS CO., LTD.,
Petitioner,

v.

NETLIST, INC.,
Patent Owner.

Case No. IPR2023-00454
U.S. Patent No. 11,093,417

DECLARATION OF STEVEN PRZYBYLSKI, PH.D. IN SUPPORT OF PATENT OWNER'S RESPONSE

TABLE OF CONTENTS

Page

I, Steven Przybylski, declare as follows:

I. INTRODUCTION

1. Netlist, Inc. has retained me to provide testimony regarding the technology involved in U.S. Patent No. 11,093,417 ("the '417 Patent") and how it compares with the asserted art.

2. I am being compensated for my work in this matter at my standard hourly rate for consulting services. My compensation is not dependent on the outcome of the case and does not affect the substance of my statements in this Declaration. I have no financial interest in the matter or the patents at issue.

3. I have reviewed the specification and the claims of the '417 Patent. I am familiar with the technology involved therein, having worked in the field for the past 33 years.

4. To prepare this Declaration, I have considered the petition, Dr. Andrew Wolfe's declaration and deposition testimony, the exhibits already submitted, my own experience and knowledge, and the documents cited herein. I provide my testimony from the perspective of a person of ordinary skill in the art at the time of the invention. I was at least a person of ordinary skill in the art at the time of the invention. I am familiar with the knowledge and skill level of these people at the time of the invention because I was intimately involved in the industry at the time and had frequent interactions with such people at the relevant

time. Unless otherwise noted, I have been asked to apply the plain and ordinary meaning of the terms from the perspective of a person of ordinary skill in the art in light of the specification and other intrinsic evidence.

II. EXPERT QUALIFICATIONS AND COMPENSATION

5. Below is a summary of my education and experience. My *curriculum vitae*, included as **Appendix A** hereto, records my education, experience, and publications in greater detail.

6. I earned a Bachelor of Applied Science from the University of Toronto in 1980. I was enrolled in the Engineering Science program, completing a course of study combining the Electrical Engineer and Computer Science options.

7. I earned a Master of Science in Electrical Engineering degree and a Ph.D. in Electrical Engineering in 1982 and 1988 respectively, both from Stanford University.

8. I also earned a Masters of Business Administration from the Haas School of Business at the University of California at Berkeley in 2000.

9. I have extensive experience with memory semiconductor integrated circuits and the memory systems constructed of them. At Stanford, my dissertation was on the optimization of single- and multi-level cache hierarchies to maximize system-level performance. Also at Stanford, I was a member of the

core team that architected, designed, built, and tested the seminal MIPS processor. I was responsible for the design of the instruction decode and control units as well as architecting the virtual memory support. I assembled and debugged the entire microprocessor design and oversaw its fabrication and testing. In 1984 and 1985, I took a leave of absence from Stanford to become a member of the founding team of MIPS Computer Systems, a startup in California that designed, built, and sold processors and computer systems. In 1989, after finishing my doctorate and brief post-doctorate at Stanford, I returned to MIPS Computer Systems. Throughout my two periods of employment at MIPS Computer Systems, I had a number of titles and worked on a variety of projects. Noteworthy were stints as processor architect, systems architect, hardware-software team liaison, and Chief Scientist of the High-End Systems group. During the second period, I also served as a Consulting Assistant Professor in the Department of Electrical Engineering at Stanford. I designed and taught an advanced graduate level course on cache and memory system design.

10. In 1991, I left MIPS to become an independent consultant. In that capacity I have provided technical, strategic, marketing, and intellectual property services to a large number of clients both domestically and internationally.

11. I am a member of both the Institute of Electrical and Electronics Engineers (IEEE) and the Association of Computing Machinery (ACM). In

addition, in 1997 I became a registered patent agent. I am listed as the inventor on 11 domestic patents and their foreign counterparts. The subject matters include memory devices, memory subsystems, networks, error detection and correction codes, processors and cache memories. Between 1994 and approximately 2000, I was a leading independent analyst following, writing, and teaching about both the technical and business aspects of the DRAM industry. Overall, I have written three books and over 25 articles in the areas of memory, processors, and computer systems. In particular, I wrote and revised an 850-page book entitled "New DRAM Technologies: A Comprehensive Analysis of the New Architectures." I have presented many lectures, classes and seminars on semiconductor memories and their memory systems, processors and computer systems. I have presented full and half-day seminars on discrete and embedded DRAMs to over 2,500 engineers, marketers, managers, and financial analysts.

III. TECHNOLOGY BACKGROUND

A. DRAMs

1. DRAM Basics

12. The term DRAM is an acronym for Dynamic Random Access Memory. A DRAM is a semiconductor memory device and an integrated circuit, with a DRAM die enclosed in a package. It is a random access memory in that, subject to certain restrictions, the storage locations, each of which store a single

bit of information, can be accessed in any order. DRAMs are dynamic in that they do not store the information written to them for very long – only a few tens of milliseconds. In order for the storage cells to retain the information for longer, the storage cells need to be periodically refreshed.

13. Logically the storage cells within a DRAM are arranged in a collection of arrays organized into rows and columns. At the intersection of each row and column is a storage cell. Figure 7.3 from Memory Systems; Cache, DRAM, Disk illustrates this logical organization.

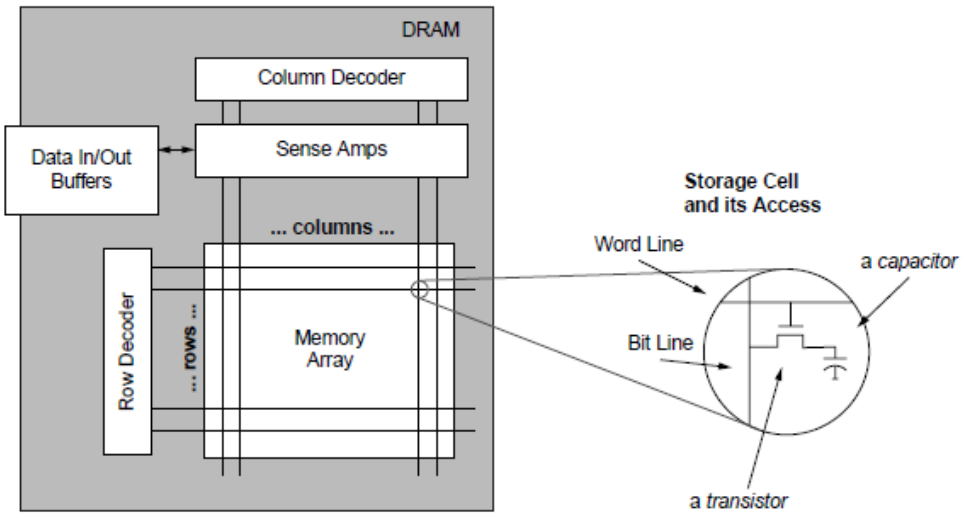

FIGURE 7.3: Basic organization of DRAM internals. The DRAM memory array is a grid of storage cells, where one bit of data is stored at each intersection of a *row* and a *column*.

Ex. 1070, p. 317.

14. The collection of cells in one logical memory array is called a bank, and a DRAM can contain a number of banks. Each of the banks is given a unique identifier or address; each row in the logical array is also given a unique row

address and each column is given a unique column address. Each storage cell is therefore uniquely identified by the aggregation of a bank address, a row address, and a column address.

15. Most DRAMs have a fixed "width" which is the number of bits that can be input or output from the DRAM simultaneously on that many discrete I/O lines. For example, a DRAM with 4 I/O lines is called a "by 4" DRAM, and this is represented symbolically with "x4". Since each collection of 4 bits is always accessed together, they are jointly assigned a single column address.

16. Thus, for example 2Gb (giga-bit) DDR4 SDRAM with 2 billion storage cells (actually 2,147,483,648 storage cells) in a 256M x 4 configuration is organized as 8 banks, each logically organized with 32,768 row, each with 4096 storage cells on it. Thus each row in a bank is identified by a 14 bit row address. Each row logically organized into 2,048 groups of 4 cells. Each group of 4 bits is input or output together, and each group is identified by a column address that is 11 bits wide. Ex. 2002 p. 7.

17. Accessing specific storage cells involves a two-step process. First the row that the storage cells are included in must be activated with an activate command that specifies the row to be accessed. This brings all the information from all of the storage cells in that row into the sense amplifiers, illustrated in the Figure 7.3 above from Ex. 1070. Once that operation is complete, then those

sense amplifiers can be read and written with read and write commands, each of which specifies a column address. Because read and write commands access specific columns within the currently activated row, they are called column access commands. Multiple reads and writes can be performed on the activated row while it is activated. Once all of the desired reads and writes to that row have been completed, the bank must be precharged before another (different) row can be activated for reading and writing. Precharging ensures that the information for the currently activated row is safely stored in storage cells for that row, the row is deactivated, and the bank is in a state in which it is ready for the next activation command.

18. Modern DRAMs are generally burst devices. That means that not only are multiple bits input and output in parallel, but each read or write command causes multiple such groupings to be input or output. For example, if an x8 DRAM is performing bursts of 4, a read to column 16 would output the 8 bits that shared that column address immediately followed by the 8 bits with column address 17, then 18, then 19. In total 32 bits of information would be output in 4 chunks of 8 bits in rapid succession all from one read column access command.

2. JEDEC-Standardized SDRAMs

19. JEDEC (the Joint Electron Device Engineering Council) is an independent standard-setting organization. It develops standards for a host of semiconductor memory and non-memory products through the cooperative efforts of its member companies. With regard to the standards it develops pertaining to DRAMs and memory modules, standards are developed through the input, critique and refinement the best and most experienced engineers from both semiconductor manufacturers and many systems companies (https://www.jedec.org/about-jedec/member-list). These JEDEC meetings are where these two bodies of knowledge are brought together to create standards that are both technically effective for use in systems and cost-effectively manufacturable by the semiconductor vendors. These standards include specifications necessary to ensure multiple-sourcing of compliant parts. These standard include the physical, mechanical, electrical, temporal and/or behavioral aspects of the devices standardized.

20. In the present matter, based on the Petition, JEDEC standards are at issue in relation to two generations of DRAMs, called DDR SDRAMs and DDR2 SDRAMs. The standard that define the interfaces and behaviors of the former are designated JESD79 with different iterations of the standard having a letter suffice. Specifically, the first DDR SDRAM standard, release in June

2000, was JESD79. Subsequent revisions were JESD79A through JESD79F, released in February 2008. For the most part, the revision did not include substantial or significant changes. The subsequent generation of SDRAMs are called DDR2 SDRAMs and the standards that define their interfaces and behavior are designated JESD79-2. The first one, JESD79-2, was released in September 2003, and the most recent was JESD79-2F, which was released in February 2008.

a) DDR2 SDRAMs

21. A POSTIA would recognize that DDR2 SDRAMs refer to certain synchronous DRAMs standardized by JEDEC. DDR2 SDRAMs are part of a larger family of synchronous DRAMs that are differentiated from the asynchronous DRAMs standardized by JEDEC. DDR2 SDRAMs and all SDRAMs are synchronous devices in that they receive a clock signal (in some cases differentially) that is used to synchronize operations and address, control and/or data transfers.

(1) Pin-level and Physical Interfaces

22. DDR2 SDRAMs all have the same pins, or connections, to the rest of the system. They are all defined in Section 1.2 of JESD72-2:

1.2 Input/Output Functional Description

Symbol	Type	Function
CK, $\overline{CK}$	Input	**Clock:** CK and $\overline{CK}$ are differential clock inputs. All address and control input signals are sampled on the crossing of the positive edge of CK and negative edge of $\overline{CK}$. Output (read) data is referenced to the crossings of CK and $\overline{CK}$ (both directions of crossing).
CKE	Input	**Clock Enable:** CKE HIGH activates, and CKE Low deactivates, internal clock signals and device input buffers and output drivers. Taking CKE Low provides Precharge Power-Down and Self Refresh operation (all banks Idle), or Active Power-Down (row Active in any bank). CKE is synchronous for power down entry and exit, and for self refresh entry. CKE is asynchronous for self refresh exit. CKE must be maintained high throughout read and write accesses. Input buffers, excluding CK, $\overline{CK}$, ODT and CKE are disabled during power-down. Input buffers, excluding CKE, are disabled during self refresh.
$\overline{CS}$	Input	**Chip Select:** All commands are masked when $\overline{CS}$ is registered HIGH. $\overline{CS}$ provides for external Rank selection on systems with multiple Ranks. $\overline{CS}$ is considered part of the command code.
ODT	Input	On Die Termination: ODT (registered HIGH) enables termination resistance internal to the DDR2 SDRAM. When enabled, ODT is only applied to each DQ, DQS, $\overline{DQS}$, RDQS, $\overline{RDQS}$, and $\overline{DM}$ signal for x4/x8 configurations. For x16 configuration ODT is applied to each DQ, UDQS/$\overline{UDQS}$, LDQS/$\overline{LDQS}$, UDM, and LDM signal. The ODT pin will be ignored if the Extended Mode Register (EMRS) is programmed to disable ODT.
$\overline{RAS}$, $\overline{CAS}$, $\overline{WE}$	Input	**Command Inputs:** $\overline{RAS}$, $\overline{CAS}$ and $\overline{WE}$ (along with $\overline{CS}$) define the command being entered.
DM (UDM), (LDM)	Input	**Input Data Mask:** DM is an input mask signal for write data. Input data is masked when DM is sampled HIGH coincident with that input data during a Write access. DM is sampled on both edges of DQS. Although DM pins are input only, the DM loading matches the DQ and DQS loading. For x8 device, the function of DM or RDQS/$\overline{RDQS}$ is enabled by EMRS command.
BA0 - BA2	Input	**Bank Address Inputs:** BA0 and BA1 for 256 and 512Mb, BA0 - BA2 define to which bank an Active, Read, Write or Precharge command is being applied. Bank address also determines if the mode register or extended mode register is to be accessed during a MRS or EMRS cycle.
A0 - A15	Input	**Address Inputs:** Provided the row address for Active commands and the column address and Auto Precharge bit for Read/Write commands to select one location out of the memory array in the respective bank. A10 is sampled during a Precharge command to determine whether the Precharge applies to one bank (A10 LOW) or all banks (A10 HIGH). If only one bank is to be precharged, the bank is selected by BA0, BA1. The address inputs also provide the op-code during Mode Register Set commands.
DQ	Input/Output	**Data Input/ Output:** Bi-directional data bus.
DQS, ($\overline{DQS}$) (UDQS), ($\overline{UDQS}$) (LDQS), ($\overline{LDQS}$) (RDQS), ($\overline{RDQS}$)	Input/Output	**Data Strobe:** output with read data, input with write data. Edge-aligned with read data, centered in write data. For the x16, LDQS corresponds to the data on DQ0-DQ7; UDQS corresponds to the data on DQ8-DQ15. For the x8, an RDQS option using DM pin can be enabled via the EMRS(1) to simplify read timing. The data strobes DQS, LDQS, UDQS, and RDQS may be used in single ended mode or paired with optional complementary signals $\overline{DQS}$, $\overline{LDQS}$, $\overline{UDQS}$, and $\overline{RDQS}$ to provide differential pair signaling to the system during both reads and writes. An EMRS(1) control bit enables or disables all complementary data strobe signals.
NC		**No Connect:** No internal electrical connection is present.
V_{DDQ}	Supply	**DQ Power Supply:** 1.8V +/- 0.1V
V_{SSQ}	Supply	**DQ Ground**
V_{DDL}	Supply	**DLL Power Supply:** 1.8V +/- 0.1V
V_{SSDL}	Supply	**DLL Ground**
V_{DD}	Supply	**Power Supply:** 1.8V +/- 0.1V
V_{SS}	Supply	**Ground**
V_{REF}	Supply	**Reference voltage**

Ex. 1064 (JESD79-2), p. 6.

23. The packages that DDR2 SDRAMs are put into are also defined and the location of each of the defined pins is specified. Any DRAMs that have fewer or additional signal or supply pins or have signals assigned to other locations in the packages would not be compliant with this standard and would not be DDR2 SDRAMs.

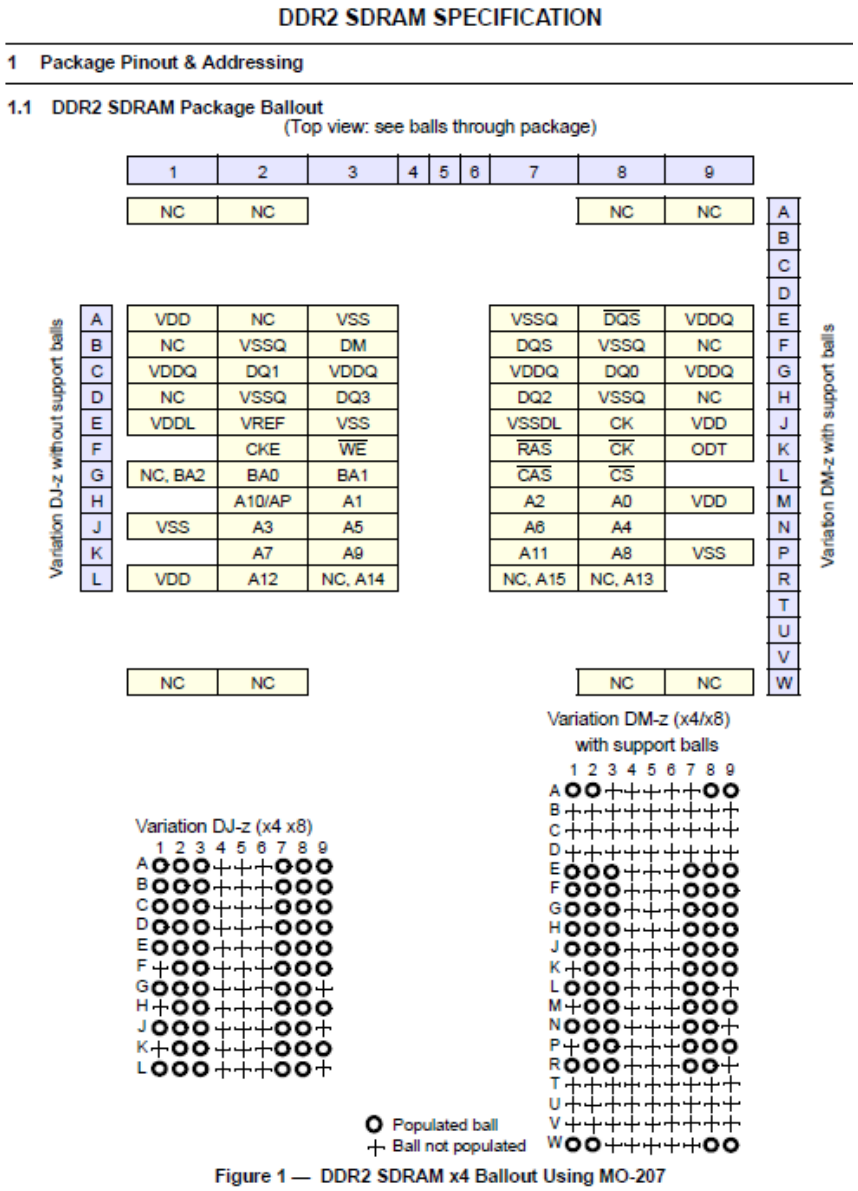

Figure 1 — DDR2 SDRAM x4 Ballout Using MO-207

Ex. 1064 (JESD79-2), p. 1.

24. Characteristic signals included in every DDR2 SDRAM are CK and $\overline{CK}$ (together called Clock in Section 1.2, above); $\overline{RAS}$, $\overline{CAS}$, and $\overline{WE}$ (called Command Inputs in Section 1.2 above); $\overline{CS}$ (Chip Select); BA0-BA2 and A0-A15 (3 bank address inputs and 16 address inputs used for row addresses and

column addresses); DQ (the bi-directional data bus, the width of which is the width of the particular SDRAM; fixed widths of 4, 8, and 16 are defined as part of the standard); and DQS (a data strobe signal that conveys timing information in parallel with the data bus). Ex. 1064 (JESD79-2), pp 1-6.

(2) Behavioral and Temporal Specification

25. DDR2 SDRAMs receive memory access commands on the command signal lines and the $\overline{CS}$. The command, along with any address (row or column) specified by the command are all captured on the rising edge of the clock signal CK.

Table 10 — Command truth table.

Function	CKE		$\overline{CS}$	$\overline{RAS}$	$\overline{CAS}$	$\overline{WE}$	BA0 BA1 BA2	A15-A11	A10	A9 - A0	Notes
	Previous Cycle	Current Cycle									
(Extended) Mode Register Set	H	H	L	L	L	L	BA	OP Code			1,2
Refresh (REF)	H	H	L	L	L	H	X	X	X	X	1
Self Refresh Entry	H	L	L	L	L	H	X	X	X	X	1
Self Refresh Exit	L	H	H	X	X	X	X	X	X	X	1,7
			L	H	H	H					
Single Bank Precharge	H	H	L	L	H	L	BA	X	L	X	1,2
Precharge all Banks	H	H	L	L	H	L	X	X	H	X	1
Bank Activate	H	H	L	L	H	H	BA	Row Address			1,2
Write	H	H	L	H	L	L	BA	Column	L	Column	1,2,3,
Write with Auto Precharge	H	H	L	H	L	L	BA	Column	H	Column	1,2,3,
Read	H	H	L	H	L	H	BA	Column	L	Column	1,2,3
Read with Auto-Precharge	H	H	L	H	L	H	BA	Column	H	Column	1,2,3
No Operation	H	X	L	H	H	H	X	X	X	X	1
Device Deselect	H	X	H	X	X	X	X	X	X	X	1
Power Down Entry	H	L	H	X	X	X	X	X	X	X	1,4
			L	H	H	H					
Power Down Exit	L	H	H	X	X	X	X	X	X	X	1,4
			L	H	H	H					

NOTE 1 All DDR2 SDRAM commands are defined by states of $\overline{CS}$, $\overline{RAS}$, $\overline{CAS}$, $\overline{WE}$ and CKE at the rising edge of the clock.

NOTE 2 Bank addresses BA0, BA1, BA2 (BA) determine which bank is to be operated upon. For (E)MRS BA selects an (Extended) Mode Register.

NOTE 3 Burst reads or writes at BL=4 cannot be terminated or interrupted. See sections "Reads Interrupted by a Read" and "Writes Interrupted by a Write" in section 2.2.4 for details.

NOTE 4 The Power Down Mode does not perform any refresh operations. The duration of Power Down is therefore limited by the refresh requirements outlined in section 2.2.7.

NOTE 5 The state of ODT does not affect the states described in this table. The ODT function is not available during Self Refresh. See section 2.2.2.4.

NOTE 6 "X" means "H or L (but a defined logic level)".

NOTE 7 Self refresh exit is asynchronous.

Ex. 1064 (JESD79-2), p. 49.

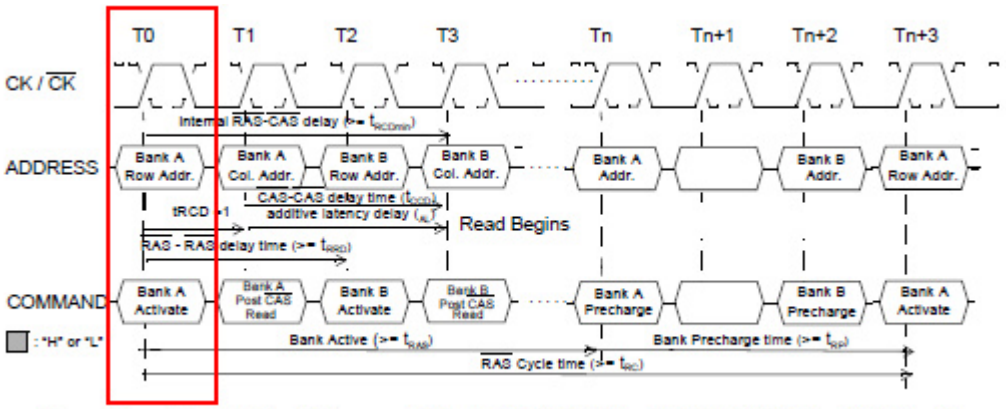

Figure 20 — Bank Activate Command Cycle: tRCD = 3, AL = 2, tRP = 3, tRRD = 2, tCCD = 2

Ex. 1064 (JESD79-2), p. 23

26. In the case of read or write commands, data is transferred later. Read data is output temporally aligned with the clock signal, but write data – data to be written into the memory array as a part of a write command – can arrive at the DRAM not aligned with the clock. The DQS data strobe signal travels with the data and arrives at the DRAM with its edged centered on each valid piece of data in a burst. The strobe DQS is thus used by the DRAM to reliably capture the data regardless of the offset of the data DQ relative to clock CK. This facility that allows data to arrive at the DRAMs at a different time than the commands and address is central to how these SDRAMs are aggregated into higher-bandwidth memory systems.

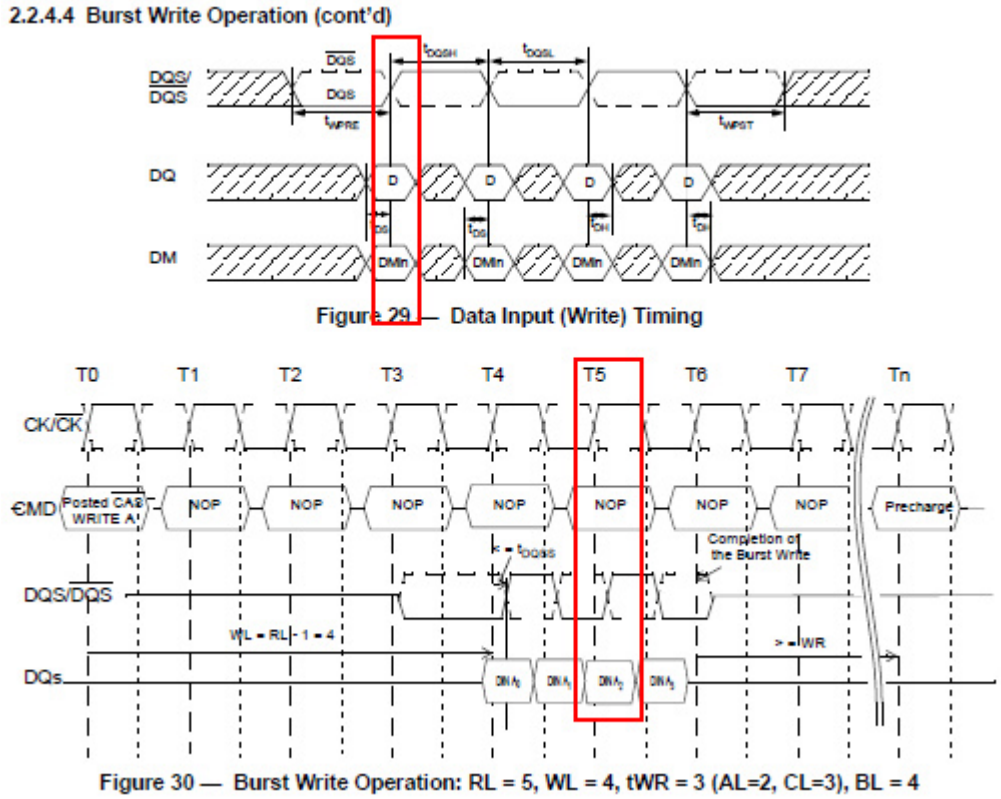

Ex. 1064 (JESD79-2), p. 30.

27. For read data being output from the SDRAM, the DQS is output with the data and can be used by the recipient of the data – the memory controller or an intermediate buffer – to help capture the data.[1]

[1] I understand that Dr. Harold Stone, Micron's expert in the district court case, has testified that "some people may consider the DQS signals … also as control signals." IPR2022-00615, EX2117, 59:2-9.

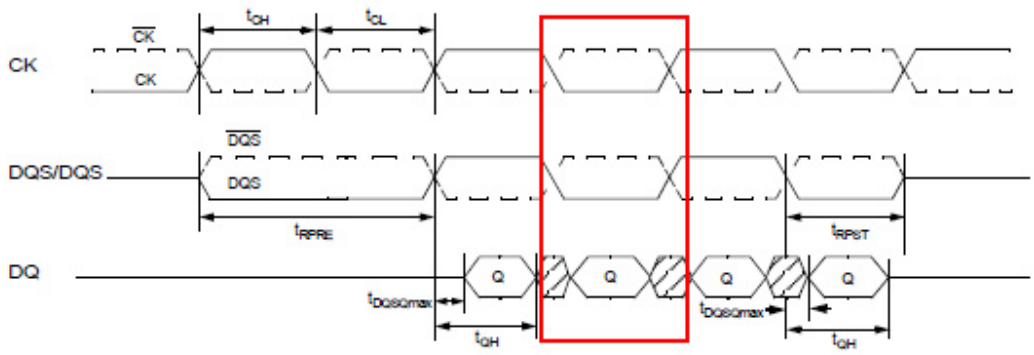

Figure 23 — Data Output (Read) Timing

Ex. 1064 (JESD79-2), p. 26.

28. So, for both read data and write data, the DQS signal travels from source to destination along with the data to provide the recovering device timing information to facilitate capturing the data. In both case, the period of the toggling of the strobe signal DQS is the same as the period of the clock CK. In both cases, one piece of data is transferred for each phase (or half period) of DQS. Thus, for each clock period during a data transfer, there are two edges of DQS and two pieces of data transferred. "DDR" stands for Double Data Rate and indicates that these SDRAMs transfer two pieces of data per clock period.

29. The time delay between the receipt of a read command and the output of the first piece of data of the burst by a memory device is called the read latency, abbreviated RL. It is measured from the rising edge of clock CK that captures the read command to the rising edge of clock CK just before the first piece of data is output. Under ordinary operating circumstances, the read latency

is equal to the value programmed into the CAS Latency ("CL") field of the mode register of the DRAM (when the additive latency ("AL") field is set to zero, otherwise, RL = CL + AL). Ex. 1064 (JESD79-2), p.24. The memory controller or its controlling software must be cognizant of the SDRAMs capabilities and program the CAS Latency field appropriately so that data will be read and written reliably at the particular frequency at which the SDRAM is being operated. Ex. 1064 (JESD79-2), p. 12, 26, 64.

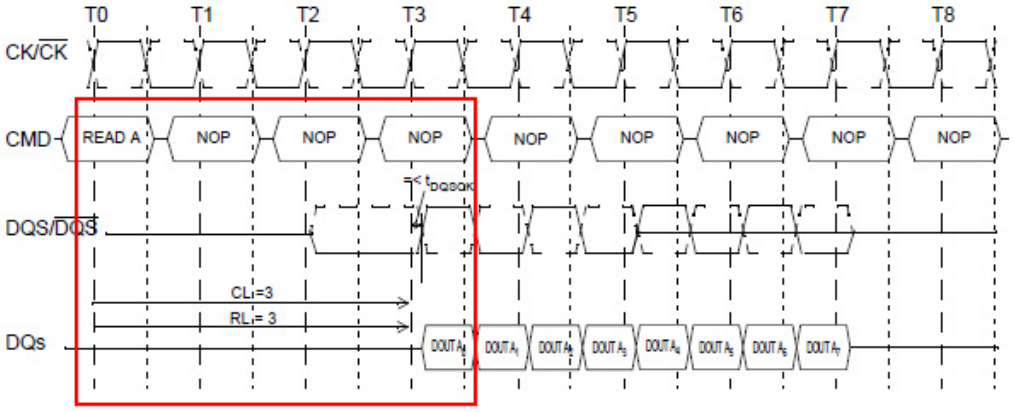

Figure 25 — Burst Read Operation: RL = 3 (AL = 0 and CL = 3, BL = 8)

Ex. 1064 (JESD79-2), p. 27.

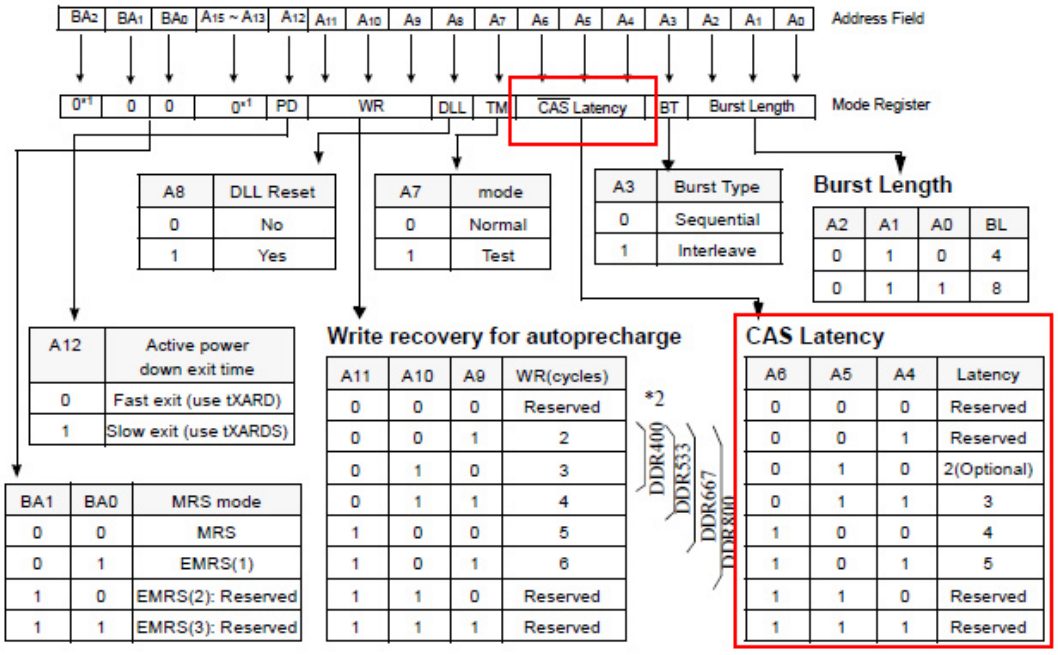

Ex. 1064 (JESD79-2), p. 12.

30. For write operations, the delay between the DRAM receiving the write command and when the DRAM expects the data to arrive is called the write latency (abbreviated WL). For DDR2 SDRAMs, the write latency is defined to be always exactly one less than the read latency (RL) and so under ordinary operating conditions is one less than the CAS Latency (CL) as programmed into the mode register. Ex. 1064 (JESD79-2), p. 12, 24, 29, 30.

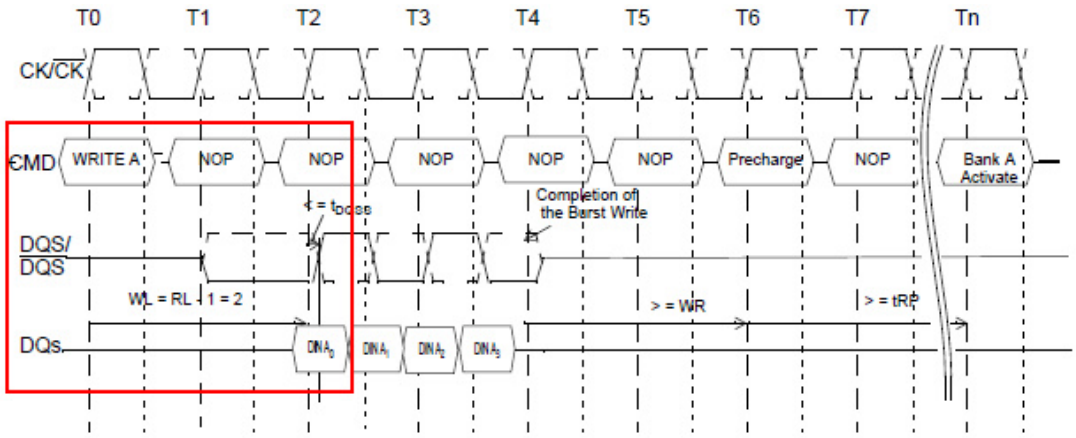

Figure 31 — Burst Write Operation: RL = 3, WL = 2, tWR = 2 (AL=0, CL=3), BL = 4

Ex. 1064 (JESD79-2), p. 30.

(3) Memory Modules

31. A memory module is typically composed of a small printed circuit board (PCB) with memory devices and support circuits mounted on one or both sides of it that is inserted into a socket on a main PCB to populate the memory system with storage. Memory modules are used in a wide variety of system types because they facilitate late selection of the memory characteristics and they are area efficient.

32. JEDEC also standardized modules using DDR SDRAM and DDR2 SDRAMs. JEDEC standardized module are called DIMMs, for Dual-Inline Memory Module, indicating that the connections between the module PCB and the connectors in the socket are on both sides of the module PCB. Ex. 1066, p. 8. The two most common styles of JEDEC-standardized modules are called Unbuffered DIMMs (or UDIMMS), which contain principally just SDRAMs,

and Registered DIMMs, which also include a register device for buffering the address and command buses (but not the data bus). Ex. 1066, p. 9.

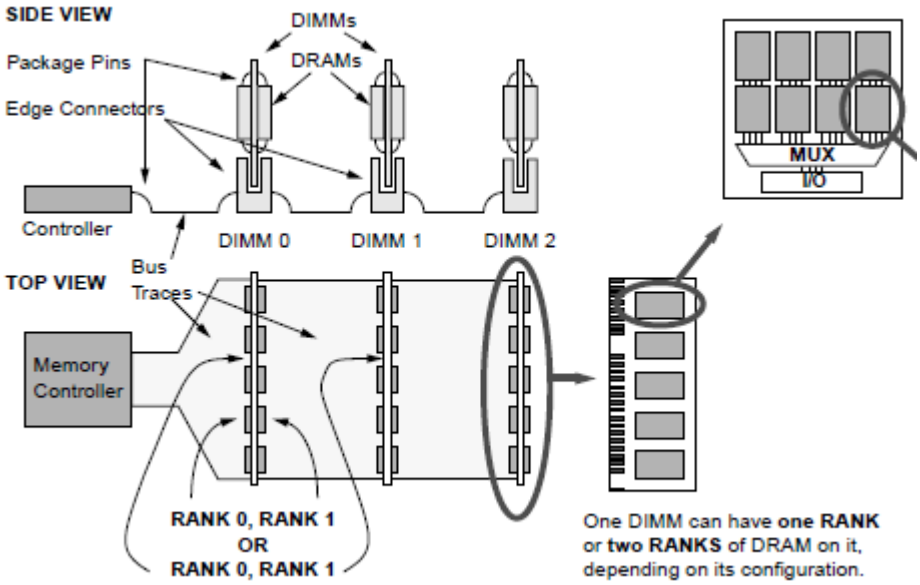

Ex. 1070, p. 319.

33. For all the principal JEDEC-standardized modules at the time of the invention, the plurality of SDRAMs are arranged logically and physically in parallel, sharing the command and address bus and presenting a data bus that is the aggregation of the width of the individual DRAMs. For example, Version A of the DDR2 RDIMM illustrated in Ex. 1066 has a 72-bit data bus at the module boundary, each of the data bits connecting to one of the DQ pins on one of the 9 x8 DDR2 SDRAMs.

Block Diagram: Raw Card Version A and F
(x72 DIMM, populated as one physical rank of x8 DDR2 SDRAMs)

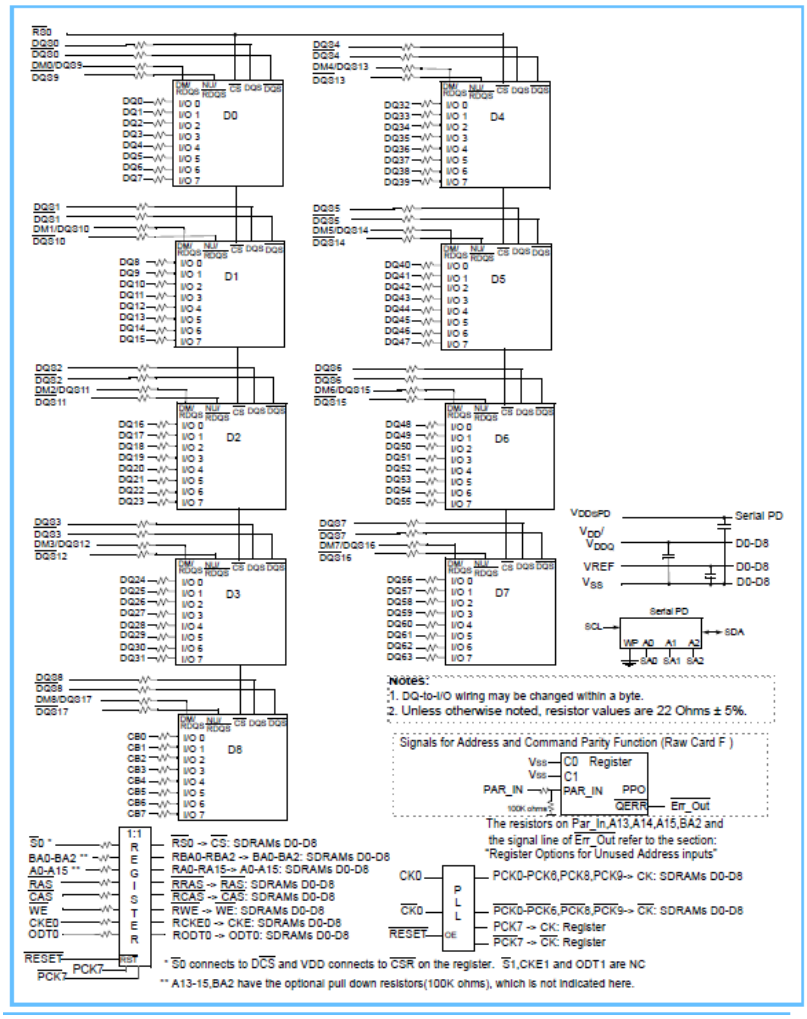

Ex. 1066, p. 9.

34. In this module, the control and address bus enters the module and is captured by the register and then distributed to each of the SDRAMs. In addition to specifying the connections to the SDRAMs in the schematics, as illustrated in the Block Diagram above, the JEDEC standard also specifies the placement of the various integrated circuits and the routing of the signals. In the case of Raw Card A, shown above, all the integrated circuits are placed on one

side of the module, and the register is placed in the middle of the module with SDRAMs placed on either side of it.

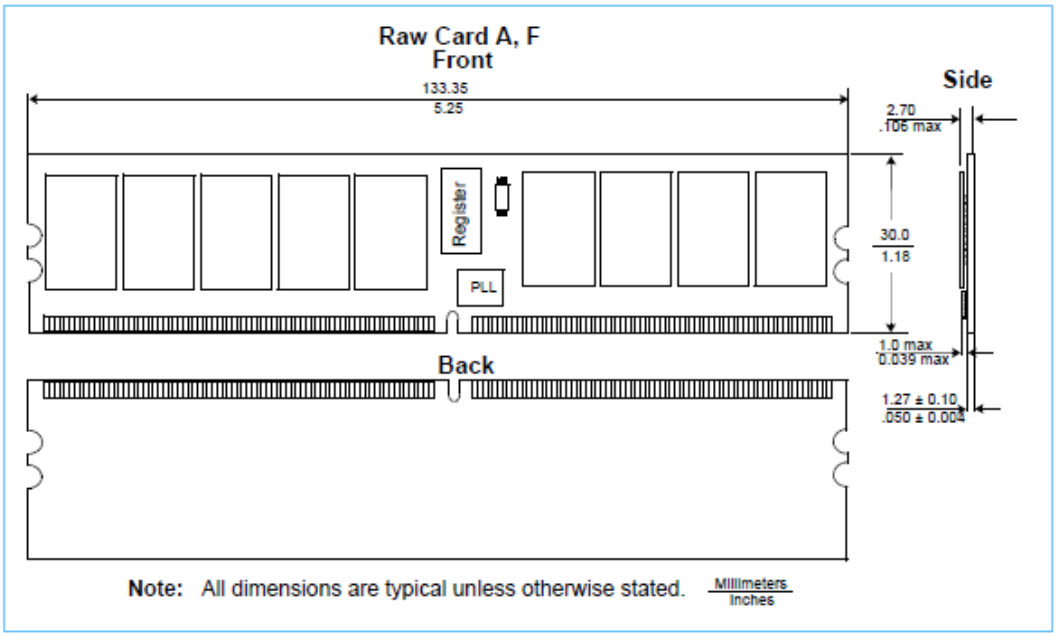

Ex. 1066, p. 44.

35. This recommended device placement is provided along with signal routing and trace length specifications that yield functional modules across the specified operating range. For example, among the many specific configurations described, the net structure routing are specified in terms of a specific topology and concrete minimum and maximum lengths of each trace segment.

Net Structure Routing for Address, Command and Control (Raw Card A, F)

A[15:0], BA[2:0], $\overline{RAS}$, $\overline{CAS}$, $\overline{WE}$, $\overline{S0}$, CKE0, ODT0

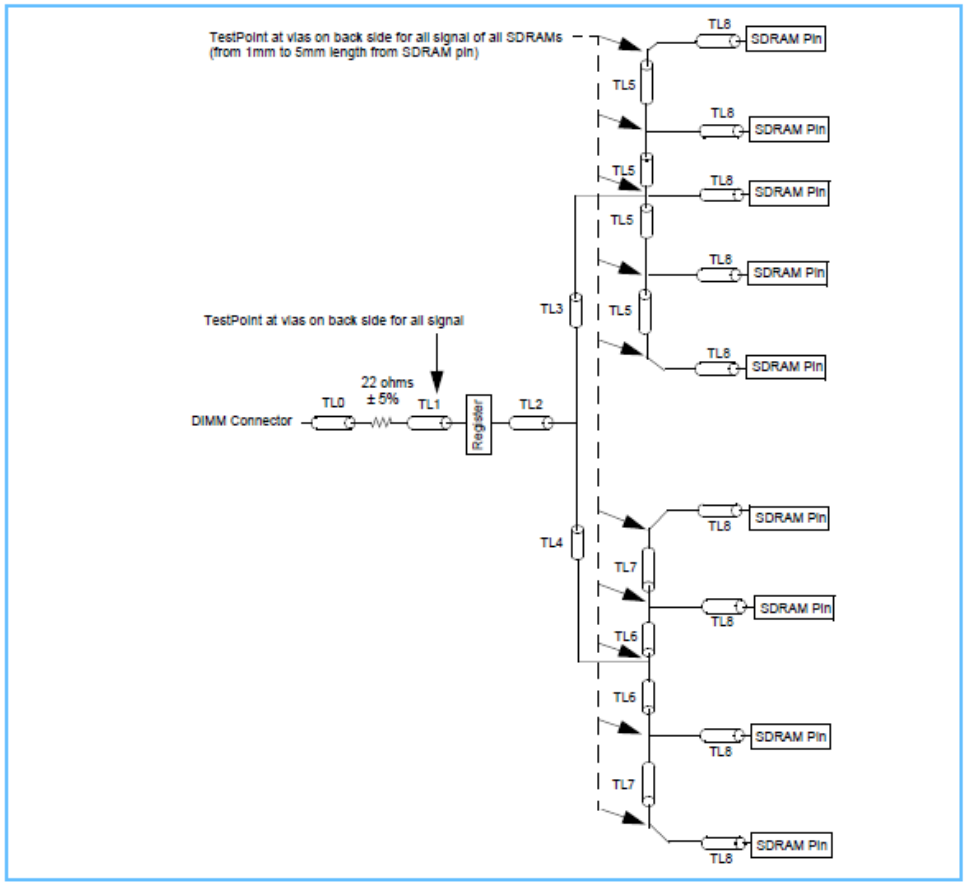

Ex. 1066, p. 65.

Trace Lengths for Address, Command and Control (Raw Card A, F)

A[15:0], BA[2:0], $\overline{RAS}$, $\overline{CAS}$, $\overline{WE}$, $\overline{S0}$, CKE0, ODT0

Raw Card	TL0		TL1		TL1 + TL2		TL2		TL3		TL4		TL5		TL6		TL7		TL8		Notes
	Min	Max	Min	Max	Min	Max	Min	Max	Min	Max	Min	Max	Min	Max	Min	Max	Min	Max	Min	Max	
A, F	3.1	4.6	28.8	30.2	33.4	33.6	0.5	0.6	37.6	37.8	46.4	46.5	13.9	14.0	6.9	7.1	13.9	14.0	1.3	5.4	1

1. All distances are given in millimeters and must be kept within a tolerance of ± 0.8 millimeter.

Ex. 1066, p. 65.

36. All of the RDIMMs specified in this standard have the register and clock regeneration PLL in the center of the module, and clock, control and address traces carefully balanced to minimize the difference in the arrival time

of these signals at each of the SDRAMs on the module. Ex. 1066, p. 4, 53, 54, 59, 65.

3. Rambus DRAMs

37. Rambus Inc. is a company that has designed and licensed a number of proprietary DRAM architectures. The best known of these architectures is called Direct RDRAM ("DRDRAMs"). DRAMs conforming to Rambus's designs, collectively called RDRAMs, do not conform to any JEDEC standard relating to the operation of DRAMs.

a) Direct Rambus DRAMs

38. Direct RDRAMs are DRAMs like DDR SDRAMs in that their storage cells are dynamic and they are logically organized into banks of rows and columns. But the interfaces by which they communicate with the memory controller are very different.

(1) Pin-level and Physical Interfaces

39. The pins of a Direct RDRAM are unique to this architecture. Instead of the separate command and address pins of the SDRAM architectures, Direct RDRAMs include a packetized 8-bit command bus RQ0..7, which is logically divided into two portions: 3 wires (RQ5..7) carry row access commands and row addresses, and 5 wires (RQ0..4) carry column access

commands and column addresses. This allows simultaneous transmission of row command packets and column command packets to the DRAMs.

Pins	Signal	I/O	Type	Description
S5, S4	SIO1 SIO0	I/O	CMOS	Serial input/output. Pins for reading from and writing to the control registers using a serial access protocol. Also used for power management.
A4	CMD	I	CMOS	Command input. Pins for reading from and writing to the control registers. Also used for power management.
A5	SCK	I	CMOS	Clock input. Clock source used for reading from and writing to the control registers
C4, D1, H1, H4, H6, K1, N1, P4	V_{DD}			Supply voltage for the RDRAM core and interface logic.
E4, M4	V_{CMOS}			Termination voltage for RSL load resistors.
A2, B4, B5, C1, C5, D4, E1, F1, F4, G4, J1, J4, J6, K4, L1, L4, M1, N4, P1, P5, R4, R5, S2	GND			Ground reference for RDRAM core and interface.
A3, B3, B2, C3, C2, D3, D2, E3, E2	DQA8..DQA0	I/O	RSL	Data byte A. Nine pins which carry a byte of read or write data between the channel and the RDRAM.
F3	CFM	I	RSL	Clock from master. Interface clock used for receiving RSL signals from the channel. Positive polarity.
F2	CFMN	I	RSL	Clock from master. Interface clock used for receiving RSL signals from the channel. Negative polarity
G1	V_{REF}			Logic threshold reference voltage for RSL signals
G2	CTMN	I	RSL	Clock to master. Interface clock used for transmitting RSL signals to the channel. Negative polarity.
G3	CTM	I	RSL	Clock to master. Interface clock used for transmitting RSL signals to the channel. Positive polarity.
H2, H3, J2	RQ7..RQ5 or ROW2..ROW0	I	RSL	Row access control. Three pins containing control and address information for row accesses.
J3, K2, K3, L2,L3	RQ4..RQ0 or COL4..COL0	I	RSL	Column access control. Five pins containing control and address information for column accesses.
S3, R3, R2, P3, P2, N3, N2, M3, M2	DQB8.. DQB0	I/O	RSL	Data byte B. Nine pins which carry a byte of read or write data between the channel and the RDRAM.

Table 1: Pin Descriptions

Ex. 2025 (Direct RDRAM Datasheet), p. 3.

40. Also, the data bus is fixed at 18 bits, though a DRDRAM can use only 16 of these if error detection/correction are not being used. There is no such thing as a x4 or x8 Direct RDRAM.

41. The clocking structure is also very different from that of SDRAMs. In those devices, there is a clock that travels with command and address signals, and the data strobes travel with data to provide timing information for its capture. Instead, with Direct RDRAMs, there are two clocks provided to the DRAM: Clock To Master (CTM) and Clock From Master (CFM), Ex. 2026 (Direct Rambus Tech. Disclosure), p. 9. The DRDRAM will receive both and use the former to capture addresses and commands on the RQ bus and write data on the DQ bus, and use the latter to read data on the DQ bus so that it can be captured by the memory controller. There is no signal that is comparable to the data strobe (DQS) or chip-select (CS) signals of SDRAMs.

42. Further, a plurality of Direct RDRAMs comprising a memory subsystem are not connected in parallel as is typical in SDRAM memory subsystems. Instead, they are all on the same linear channel. Ex. 2026 (Direct Rambus Tech. Disclosure) p. 8.

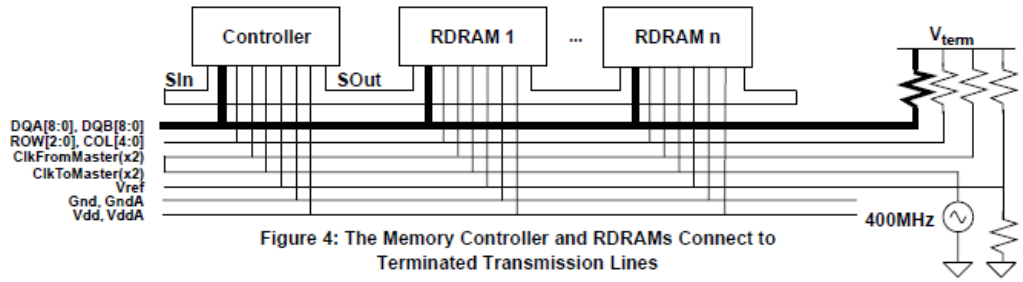

Figure 4: The Memory Controller and RDRAMs Connect to Terminated Transmission Lines

Ex. 2026 (Direct Rambus Tech. Disclosure), p. 8.

43. The packages are also different than used with SDRAMs. All Direct RDRAMs are in a chip-scale package ("CSP") that have solder balls in different number and arrangement than the packages used by SDRAMs:

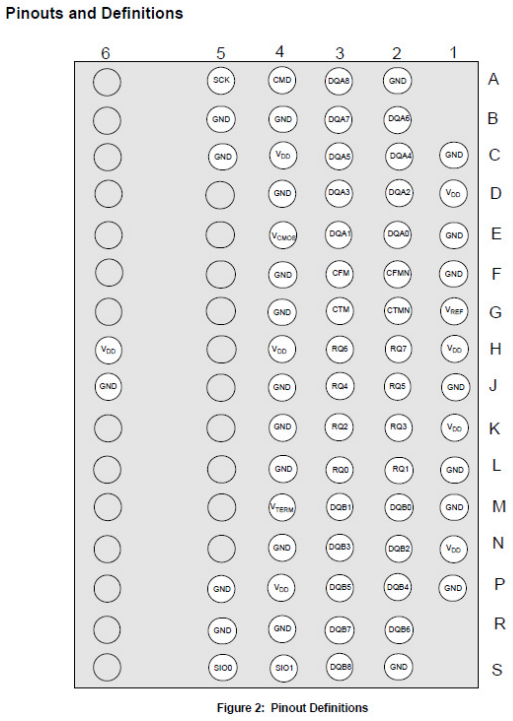

Figure 2: Pinout Definitions

Ex. 2025 (Direct RDRAM Datasheet), p. 2.

44. The high bandwidth performance of Rambus RDRAM memory systems stems directly from the linear channel topology and the attention to signal integrity. In contrast to memory systems built with JEDEC-standardized SDRAMs of all flavors (which have signal traces of several different topologies, lengths and loading) the Rambus channels are uniform, carefully terminated and using the most advanced signaling technologies, all of which contribute to much high channel frequencies than the contemporaneous SDRAM memory systems.

Ex. 2001, pp. 18 ("There are several classes of system nets in a conventions memory system"), 19, 20 ("the Direct Rambus DRAM (RDRAM) takes another approach that provides for 1.6 Gigabytes/s bandwidth"), 21.

(2) Behavioral and Temporal Specification

45. One of the major differences between Direct RDRAMs and JEDEC-standardized SDRAMs is the emphasis on bandwidth and the corresponding frequency of operation of the data and control buses. DDR SDRAMs, introduced in 2000, envision a maximum control and data bus frequency of 143 MHz, or equivalently a minimum bit time of 3.5 ns. In contrast, in Direct RDRAMs, introduced 3 years earlier, when Direct RDRAM was introduced, the data transfer frequency was 800 MHz, and the bit time on the data bus 1.25ns. The bandwidth available from a Direct RDRAM was 1.6 GBps, compared with 286MBps for x8 DDR-266A SDRAM, which is smaller by a factor of 2.8. Ex. 1060 (JESD 79), p. 10; Ex. 2025 (Direct RDRAM Data Sheet), p. 42. These higher operating frequencies were made possible by rigorous constraints on the channel physical parameters, the dedicated clocks data flowing to and from the DRAMs, and a novel low signal swing signaling architecture. Ex. 2025 (Direct RDRAM Data Sheet), pp. 1, 10, 42.

46. Another of the essential aspects of Direct RDRAMs that differentiate them from all of the JEDEC-defined SDRAMs is the encoding of

the 8-bit RQ control bus. Instead of presenting the entire command and address at once on 24 signal lines, captured by a single clock edge, the Direct RDRAM RQ bus has only 8 transmission lines. In addition, instead of 1 command spread across 24 signal lines, row commands with their row addresses are gathered into a packet and sent over 8 bit periods on only 3 lines. Column commands are sent with the column addresses, data masking and precharge indicators over 5 metal traces in the same 8 bit period. Ex. 2025 (Direct RDRAM Data Sheet), pp. 6-9

47. Finally, the addressing paradigm is different for Direct RDRAMs than for any of the JEDEC-Standardized SDRAMs. For JEDEC-Standardized SDRAM devices, the chip select signal, CS, which is part of the command bus, is used to identify which devices are going to respond to a command. Direct RDRAMs do not have a chip select signal. Instead, each device on the channel is given a unique identifier during initialization. All command packets include a 5 bit device identifier to identify which singular DRAM on the channel is to respond to the command. As Jacob puts it, the information that determines which devices will respond to a command—what he dubs the chip-select information—"travel(s) over the same set of electrical wires but at different times." Direct Rambus channels do not use dedicated address, control, data and chip-select signals. Ex. 1069, p. 7. Therefore, there are no chip-select signals as the term would be understood by a POSITA in either the Direct RDRAM or the XDR

RDRAM architectures. The passages quoted by the Board from Ex. 1069 (Paper 11, at 14) say that the original RDRAM includes chip-select information. This is not accurate. The original RDRAM being described in these paragraphs included only an address field in the command packets and did not have any dedicated fields that a POSITA would consider as chip-select information or chip-select signals as shown below.

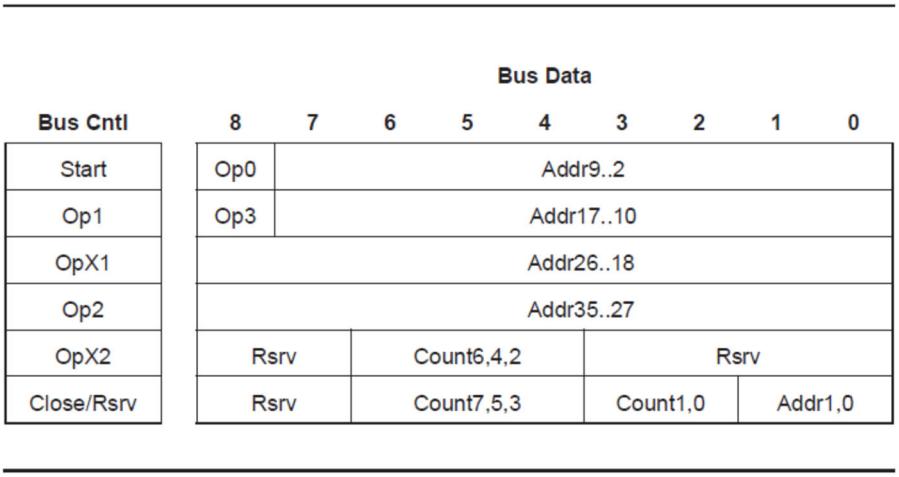

	Bus Data								
Bus Cntl	8	7	6	5	4	3	2	1	0
Start	Op0	Addr9..2							
Op1	Op3	Addr17..10							
OpX1	Addr26..18								
Op2	Addr35..27								
OpX2	Rsrv		Count6,4,2			Rsrv			
Close/Rsrv	Rsrv		Count7,5,3			Count1,0		Addr1,0	

Figure 8-8 A Rambus Request Packet takes 12 ns on the bus.

Ex. 2029, p. 277.

48. Also, the 64Mbit Direct RDRAMs, which were developed in 1997 included 16 banks. Ex. 2026 (Direct Rambus Tech Disclosure), p. 13. No other 64Mbit DRAM defined by JEDEC has that many banks. In particular, the contemporaneous 64Mb DDR SDRAM had only 4 banks. Ex. 1060, p. 8. Further, since each of the up to 32 Direct RDRAMs on a channel has 16 banks, the total number of banks within a Direct RDRAM memory systems is very

much larger than is found in a typical DDR SDRAM or DDR2 SDRAM memory system. Because there are so many banks, each bank is relatively small. Unlike the banks in the 64Mb DDR SDRAM, which have 12 row address bits and 7, 8, or 9 column address bits, (depending on the device width), the 64 Mb Direct RDRAM banks each have only 9 address bits and 6 column addresses bits. Ex. 1060, p 14; Ex. 2025 (Direct RDRAM Data Sheet), p. 7.

(3) Rambus Direct RDRAM Modules

49. The modules defined by Rambus Inc. to provide the benefits of modules to the Direct RDRAM ecosystem (called RIMMs) are very different from the unbuffered and registered DDR2 DIMMs described above. Though the size of the module is similar to that of DIMMs and the same socket is used, that is the extent of the similarity. As noted above, Direct Rambus RDRAMs are placed inline on a single channel carrying all control, address and data information. For the purpose of a module, this channel enters the module on one side, traverses the module, connecting to any Direct RDRAMs on the module and then leaves the module at the other end. The modules can have up to 16 Direct RDRAMs on them – all on the one channel. Ex. 2027 (RIMM Specification), pp. 1-1, 2-9, 2-10.

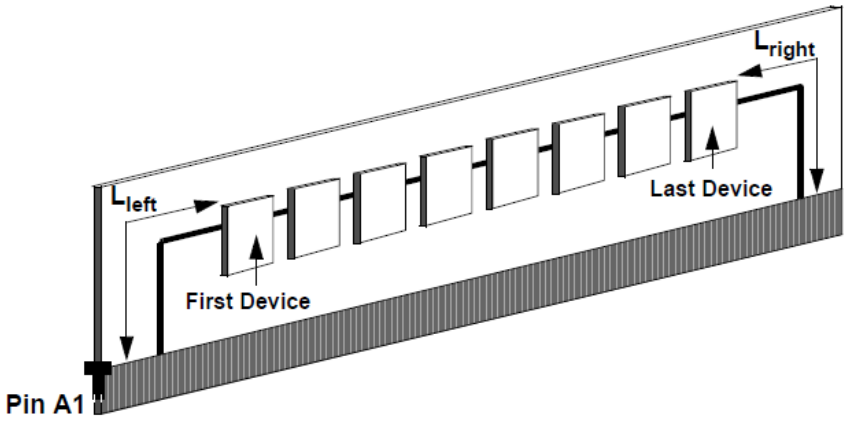

Ex. 2027 (RIMM Specification), p. 2-10.

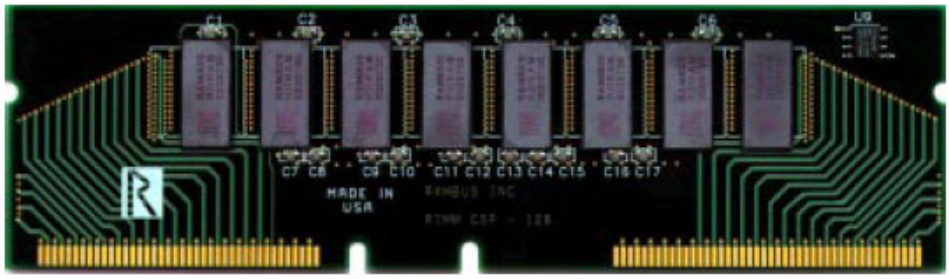

Figure 10: Physical Organization of a Rambus-Based System

Ex. 2026 (Direct Rambus Tech. Disclosure), p. 15.

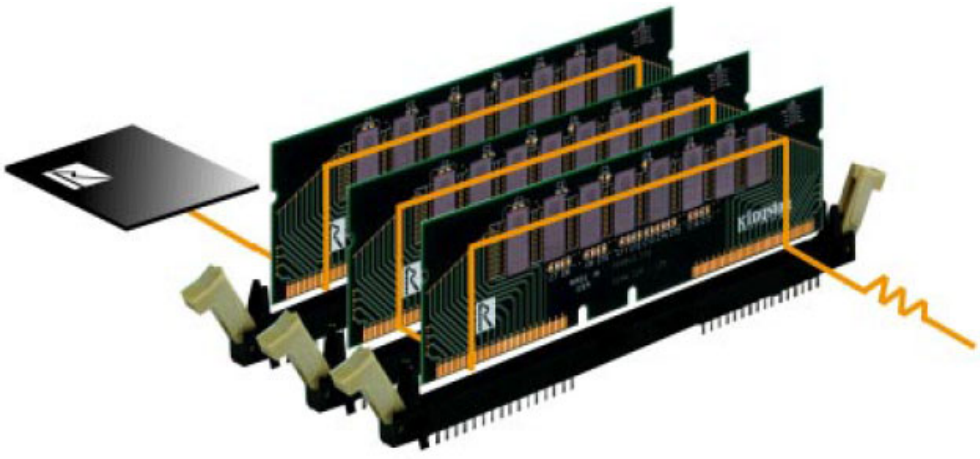

Figure 1: An Example Direct Rambus Memory Subsystem

Ex. 2026 (Direct Rambus Tech. Disclosure), p. 3.

50. Rambus modules (RIMMs) can also have zero RDRAMs on them. These "continuity modules" connect the channel to the termination resistors at the far end of the channel from the controller.

51. In JEDEC-style DIMMs, several DRAMs are "ganged together in parallel to provide the necessary aggregate bandwidth," thus increasing the size of independent banks but not their number. *See* Ex. 2001, p. 23 (Fig. 8(a), reproduced below). In contrast, "a Direct RDRAM spans the entire channel" and "the CPU accesses each RDRAM independently." *Id.*, p. 23. "So each DRAM directly adds to the number of memory banks accessible to the memory controller." *Id.*; *see* Fig. 8(b) below.

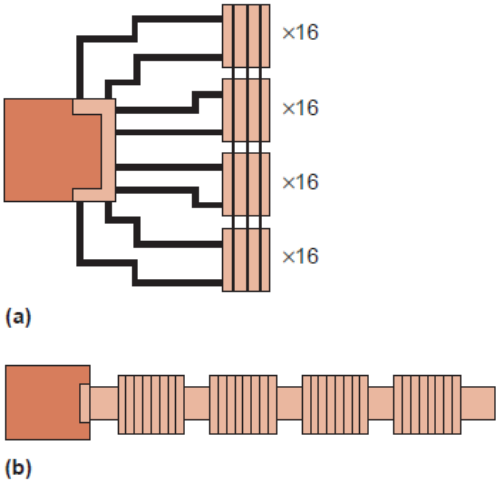

Figure 8. Bank counts: a 32-Mbyte, 64M SDRAM system with four large banks (a) versus a 32-Mbyte, 64M Direct RDRAM system with 32 small banks (b).

b) XDR Rambus DRAMs

52. After Direct RDRAMSs, Rambus also developed the XDR architecture, including XDR DRAMs and corresponding memory modules featuring a differential clock input (clock from master), request/command bus RQ11…0 and bidirectional differential data bus DQ15…0/DQN15…0. The data buses are connected point-to-point to the memory controller or buffer, but the request bus may be connected to several memory chips in parallel (multi-drop). *See*, *e.g.*, Ex. 2009, pp. 8-10.

53. XDR RDRAMs are also unique across the pantheon of memory devices in that the used portion of the native data bus width is programmable in the system. This feature works in conjunction with the point-to-point topology of the data buses to facilitate the highest performance level even in the face of changing memory system capacity. "In the XDR architecture, as the number of memory devices per request channel is increased, device width is reduced to maintain a fixed data path width." Ex. 2009 (XDR Arch), p. 11. The result is what Rambus dubbed a "dynamic point-to-point interface". Ex. 2009, pp. 10, 11, 12, 15, 17.

54. XDR DRAMs use a different package and pin configuration than JEDEC-compliant SDRAMs:

Table 1: 104ball XDR DRAM Package(Top View)

16	DQ10	DQ0	SDO	V_{DD}	GND		V_{DD}	V_{DD}	GND	DQ1	DQ11
15	DQN10	DQN0	RST						SCK	DQN1	DQN11
14	DQ6	DQ12	GND	RQ2	RQ5		RQ6	RQ8	CMD	DQ13	DQ7
13	DQN6	DQN12	V_{DD}	RQ1	V_{REF}		RQ7	RQ9	V_{DD}	DQN13	DQN7
12	V_{DD}	GND	GND			V_{DD}			GND	GND	V_{DD}
11	GND	V_{TERM}		V_{DD}	GND	GND	V_{DD}	GND		V_{TERM}	GND
10											
9											
8											
7											
6	GND	GND		V_{DD}	V_{DD}	GND	GND	GND		GND	GND
5	V_{DD}	V_{DD}	V_{TERM}			V_{DD}			V_{TERM}	V_{DD}	V_{DD}
4	DQ14	DQ4	GND	RQ3	RSRV		CFMN	RQ11	GND	DQ5	DQ15
3	DQN14	DQN4	RQ0	RQ4	RSRV		CFM	RQ10	V_{DD}	DQN5	DQN15
2	DQ2	DQ8	GND						V_{DD}	DQ9	DQ3
1	DQN2	DQN8	SDI	V_{DD}	GND		V_{DD}	GND	V_{DD}	DQN9	DQN3
Top View	A	B	C	D	E	F	G	H	J	K	L

ROW

COL

Ex. 2028 (Samsung XDR Datasheet), p. 4.

55. The commands bus RQ is packetized and each pair of 12-bit time slots can hold a single command. The 5 basic packet formats are shown in Figure 1 of the XDR Architecture specification.

Figure 1 **Request Packet Formats for Full XDR Architectural Range**

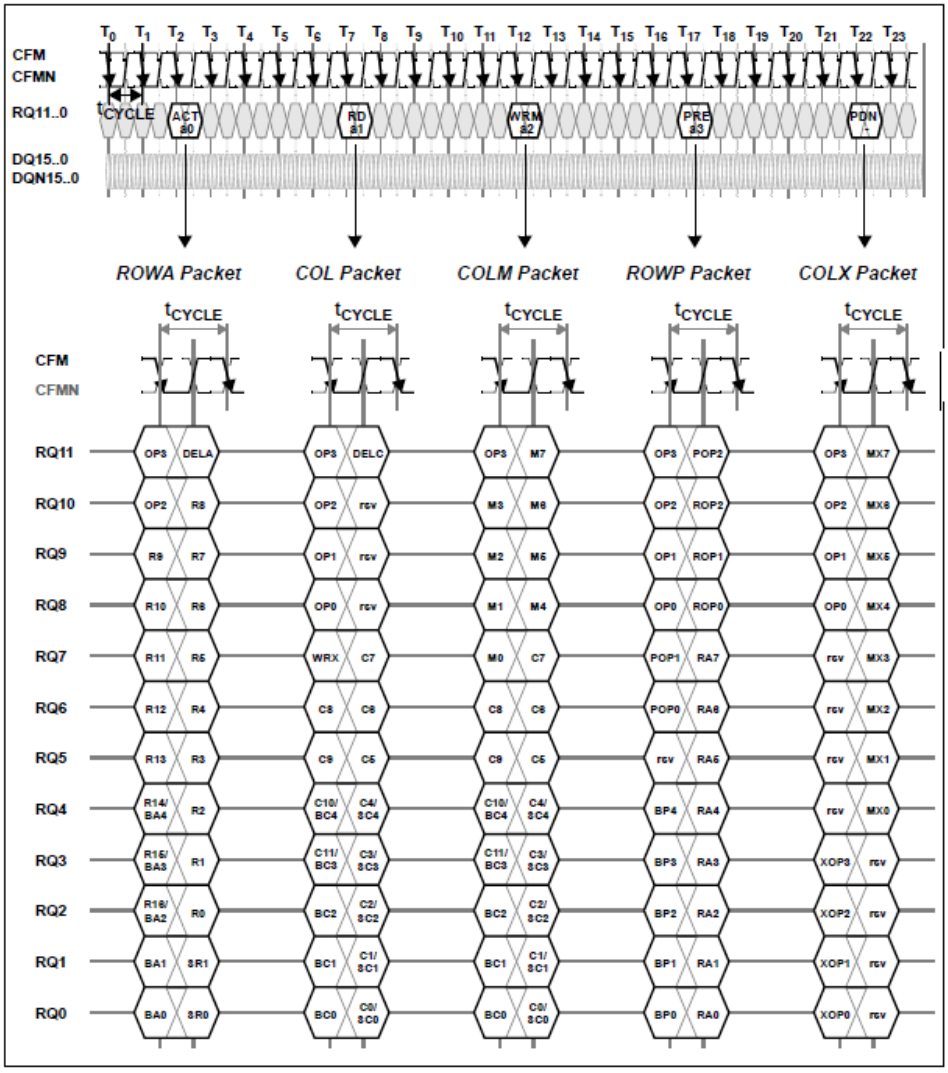

Ex. 2009 (XDR Arch.), p. 5.

56. The opcode field of 3 or 4 bits, depending on the packet type, specifies the command to be completed. There are no chip select signals or command packet fields and no device ID bits because the topology of the memory system makes these unnecessary.

57. The variable width data bus makes addressing very complex. The width of the row and column address fields in the request packets depends not only on the native device width but the programmed device width and other features such sub-column addressing, the device page size and the number of banks in the memory device. "XDR DRAMs are defined as having a native width, which is the maximum width of the DQ interface supported, and an optional range of narrower programmed widths. The following sections describe the addressing considerations as a function of the native and programmed device widths." Ex. 2009 (XDR Arch.), p. 6. "Column addressability, in turn, must proportionally scale with the number of devices in the system and, therefore, inversely with the programmed device width. The incremental column address bits used to achieve this increased addressability are used within the DRAM to select the appropriate target subset of the column for the current transaction and the data path routing to/from that section." Ex. 2009 (XDR Arch), p. 6.

IV. CLAIM CONSTRUCTION

A. "Rank"

58. In the DDR2 SDRAM RDIMMs described in Ex. 1066, certain collections of the DRAMs are referred to as ranks. This standard defines some modules with one rank, others with two ranks and still others with 4 ranks. In these modules, a "physical rank" is the collection of SDRAMs that share all the

command and address input signals, including most significantly the chip-select signal that make up part of the command. For example in the Raw Card Versions A and F, illustrated below, all of the DRAMs together as illustrated constitute a rank. Ex. 1066, p. 9. The DRAMs of Raw Card Versions B and G are arranged into two ranks. All of the DRAMs in both ranks share the same command and address signals except for the registered versions of the rank select signals S0 and S1. This collection of DRAMs that respond to commands together and provide data to the full width of a JEDEC-standardized DIMM is the canonical example of a rank and would be the example that a POSITA would give in response to the question "what is a rank".

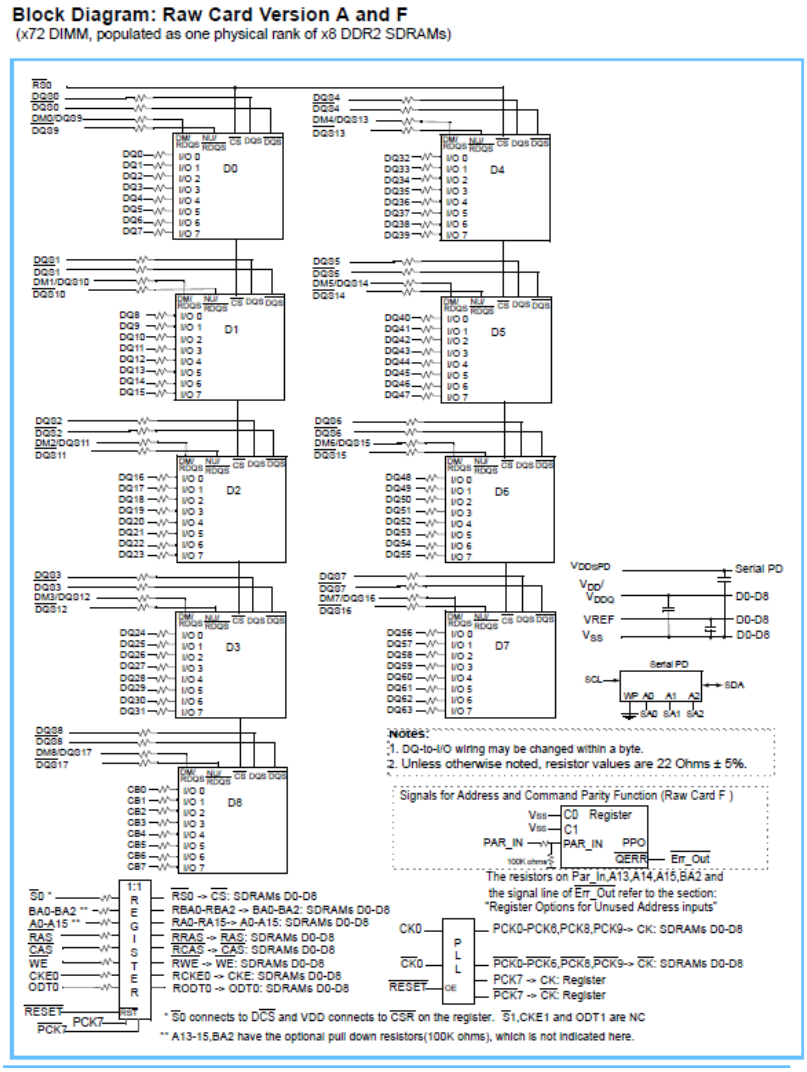

Ex. 1066, p. 9.

59. Historically, the invention of the term rank in 1995 was a direct result of the first DRAMs that were internally partitioned into banks. The overloading of the term bank to refer to both memory-system banks with the structures internal to the memory system had the potential create ambiguities. The term rank was explicitly defined in terms of the granularity of the memory system: the collection devices that make up a rank of memory has to be all

present for the memory system to operate properly. As such, the term was not used to describe single memory devices in the newest (at the time) architectures (such as Rambus) that used a narrow high-speed channel to communicate with the DRAMs individually. Instead the term was only used in the context of memory systems constructed of the asynchronous DRAMs or SDRAM and SDRAM-like memory devices.

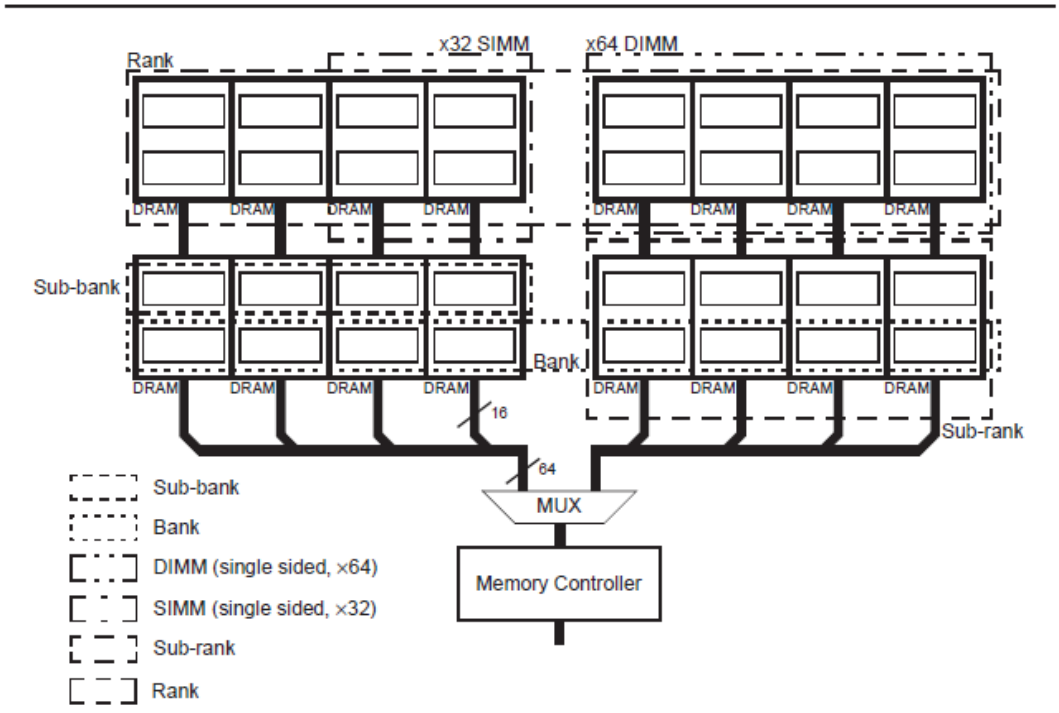

Figure 4-6 Complex memory systems are divided logically into banks and physically into ranks. DIMMs and SIMMs are ganged in parallel to form ranks, each of which may consist of one or more banks.

Ex. 2029 (New DRAM Technologies 2d Ed.), p. 79.

60. In the corresponding district court case, the parties here proposed different constructions for "rank" in the context of the '417 patent. Petitioner proposed that a "rank" is "an independent set of one or more memory devices on

a memory module that act together in response to command signals, including chip-select signals, to read or write the full bit-width of the memory module." Ex. 2030 (Claim Construction Order), p. 12. Petitioner proposes the same construction in this proceeding. Petition, p. 26.

61. Netlist, in contrast proposed that a rank is a predetermined group of memory devices on a memory module that can send or receive a fixed number of data bits via a fixed width data bus, in response to a read or write command and independently from other memory devices on the memory module. Ex. 2030 (Claim Construction Order), p. 12.

62. The district court, in its Claim Construction Memorandum Opinion and Order, issued November 21, 2023, decided that the proper construction of the term in the context of the patents before it in that proceeding is "a 'bank' of one or more devices on a memory module that operate in response to a given signal." Though this construction resolves some of the tension between it two proposed constructions, it does so by defining a rank in terms of a bank of memory devices without identifying what is and what is not a bank of memory devices. Ex. 2030 (Claim Construction Order), pp. 15, 35.

63. In particular, since the mid-1990s when the term rank was introduced into the parlance of memory systems, a reference to a bank of DRAMs has been archaic and depreciated, but still occasionally, and often

inaccurately, used by practitioners. In the pre-1994 world of small, narrow, slow and asynchronous DRAMs, a memory bank constructed DRAMs was a one or two dimensional array of devices that were aggregated together to create a wide memory subsystem. In many but not all legitimate uses of the term bank, particularly in memory systems relying of interleaving to develop higher bandwidth, the DRAMs within a bank were not homogeneously controlled. In this past era and context, a bank of a single DRAM was nonsensical and outside what a POSITA would refer to as a bank.

64. With regard to the district court construction's reference to operation in response to a signal, to the extent "a signal" means a single signal as opposed to one or more signals, not since the asynchronous DRAMs of the pre-SDRAM era have DRAMs responded to single signal on a single signal trace. In all of the memories in common use in 2004-2005, it takes a command comprised a plurality of signals across the signal traces of the command bus to make the DRAM "operate" i.e. to perform a meaningful command.

65. The court's construction is also not consistent with the plain and ordinary meaning as captured by the section 4.1.2.10 of JESD21-C, which is the standard that defines the SPD (Serial Presence Detect) for DDR2 SDRAMs. Byte 5 of the SPD is "the number of DIMM ranks", which is defined specifically in terms of the chip select pins, not any arbitrary signal.

Byte 5: Module Attributes - Number of Ranks, Package and Height

This field describes the number of ranks (Rank: any DRAMs connected to same physical $\overline{CS}$) and package on the SDRAM module, and module height. The number of logical banks for the SDRAM device is defined in Byte 17.

Bit 7 ~ Bit 5	Bit 4	Bit 3	Bit 2 ~ Bit 0
Module Height	DRAM Packgage	Card on Card	# of Ranks
Bit[7, 6, 5] 000 = less than 25.4mm 001 = 25.4mm 010 = greater than 25.4mm and less than 30mm 011 = 30.0mm 100 = 30.5mm 101 = greater than 30.5 mm others :reserved	1 = stack 0 = planar	1 = yes 0 = no	Bit [2, 1, 0] : 000 = 1rank 001 = 2ranks 010 = 3ranks 011 = 4 ranks ..... 111 = 8ranks

Ex. 2031 (JESD21-C 4.1.2.10), p. 7.

V. PERSON OF ORDINARY SKILL IN THE ART

66. I have been asked to consider Samsung's definition of a person of ordinary skill in the art for this declaration, in particular the supposed knowledge that such persons would have had at the time of the invention. In his Declaration, Dr. Wolfe states:

> I believe a person of ordinary skill in the art in the field of the 417 Patent in 2004 or 2005 ("Skilled Artisan" or "POSITA") would have been someone with an advanced degree in electrical or computer engineering and at least two years working in the field, or a bachelor's degree in such engineering disciplines and at least three years working in the field. Additional training can substitute for educational or research experience, and vice versa. Such a hypothetical person would have been familiar with various standards of the day including the JEDEC industry standards, and knowledgeable about the design and operation of standardized DRAM and SDRAM memory devices and memory modules and how they interacted with the memory controller of a computer system. Specifically, he or she would have been knowledgeable about the JEDEC DDR and DDR2 SDRAM standards used to standardize the functioning of memory devices, and the JEDEC 21-C standard used to standardize different possibilities for the physical layout of memory devices on a module as well as different possibilities for density and organization of the memory

> devices to achieve a given memory capacity. He or she would also have been familiar with the structure and operation of circuitry used to access and control computer memories, including sophisticated circuits such as ASICs (Application-Specific Integrated Circuits) and CPLDs (Complex Programmable Logic Devices) and more low-level circuits such as tri-state buffers, flip flops and registers.

Ex. 1003, ¶ 48.

67. I understand that Samsung asserts that such people would be familiar with the "various standards of the day including the JEDEC industry standards, and knowledgeable about the design and operation of standardized DRAM and SDRAM memory devices and memory modules and how they interacted with the memory controller of a computer system." Petition, p. 9. I understand that Samsung also asserts that a person of ordinary skill in the art would have been "familiar with the structure and operation of circuitry used to access and control computer memories, including sophisticated circuitries such as ASICs (Application-Specific Integrated Circuits) and CPLDs (Complex Programmable Logic Devices) and more low-level circuits such as tri-state buffers, flip flops and registers." Petition, p. 10.

68. Both the Petition and Dr. Wolfe assert that a POSITA here would be "knowledgeable about the design and operation of standardized DRAM and SDRAM memory devices and memory modules and how they interacted with the memory controller of a computer system." As noted in the specification of

the ’417 patent, “[t]he present invention relates generally to memory modules of a computer system, and more specifically to devices and methods for improving the performance, the memory capacity, or both, of memory modules.” Ex1001, 1:47-50. The truth of the matter is the design of memory modules using semiconductor memories and the design of those semiconductor memories are two distinct areas of knowledge and a POSITA with an ordinary level of skill in the design of memory modules would not possess the training, knowledge or expertise to attempt to design semiconductor memory design. DRAMs in particular are notoriously complex and subtle designs that take many years to master. Thus I believe a POSITA here would be knowledgeable about the operation of standardized DRAM and SDRAM memory devices and memory modules including these commercially available devices and how these devices and modules interacted with memory controllers of a computer system.

69. In support of my assertion here I note that the ’417 patent does not discuss the internals of DRAM semiconductor devices at all. Further, the text *Memory Systems; Cache, DRAM, Disk*, by Jacob, Ng and Wang include 285 pages about the design and operation of memory systems using DRAMs, of which only 2 pages and 1 figure describe the internal operation and structure of a DRAM in the most superficial manner. Ex. 1070, pp. 316-317. In contrast, the text “DRAM Circuit Design” by Keeth, Baker, Johnson and Lin comprise 440

pages about the design of DRAM with very little about how these devices are aggregated to form memory systems. Ex. 2032 (Keeth, et al).

70. With this restriction to the scope of knowledge and experience of a POSITA, I apply Samsung's proposed level of knowledge in my analysis.

71. The important consequence of this understanding of the knowledge and capabilities of a POSITA in these matters, is that when looking at a new memory system structure with which the POSITA is unfamiliar involving a semiconductor memory device, a POSITA would try to understand which of the commercially available memory devices that he or she is familiar with were being used in the design or whether an otherwise unknown or unfamiliar memory device is being employed. Further, confronted with a new memory module organization, again one that the POSITA is unfamiliar with, the POSITA would look to the existing modules and memory system structures to form the context for understanding the structure and operation of the unfamiliar device.

72. I note that in the 2004-2005 timeframe, the vast majority of the DRAM devices and modules sold would have been compliant with one of the JEDEC standards, such that a general reference to DRAM, especially in the context of usage main memory of general purpose computer and server systems, would suggest to a POSITA the use of JEDEC-standard compliant SDRAM devices or modules. For instance, a person of ordinary skill in the art would

recognize the "64Mb×8," "128Mb×8," "128Mb×4" DRAM devices mentioned in the '417 Patent, 2:34-46, as referring to JEDEC-standard compliant DDR2 SDRAM devices. Ex. 1060, p.7. Similarly, they would recognize the x64 and x72-wide memory modules mentioned in the '417 Patent, 2:47-53, as referencing JEDEC-standard compliant memory modules, each using the one type of JEDEC-standard compliant SDRAMs. Ex. 1066, p. 4.20.10-4 (x64 and x72 memory module organizations).

VI. THE '417 PATENT

A. General Overview

73. The '417 Patent, titled "Memory Module With Data Buffering," relates to a memory module that effects "registered transfers of N-bit wide data signals with [a] memory read or write command between [a] N-bit wide memory bus and [] memory devices in response to [] data buffer control signals and in accordance with an overall CAS latency of the memory module, which is greater than an actual operational CAS latency of the memory devices." EX1001, Abstract. The data buffer control signals are output by a logic on the module, which is "configurable to receive a set of input address and control signals associated with [the] read or write memory command and output registered address and control signals [in addition to] data buffer control signals." *Id.*

74. Figure 1, reproduced below with annotation, illustrates such a memory module. EX1001, 5:65-67

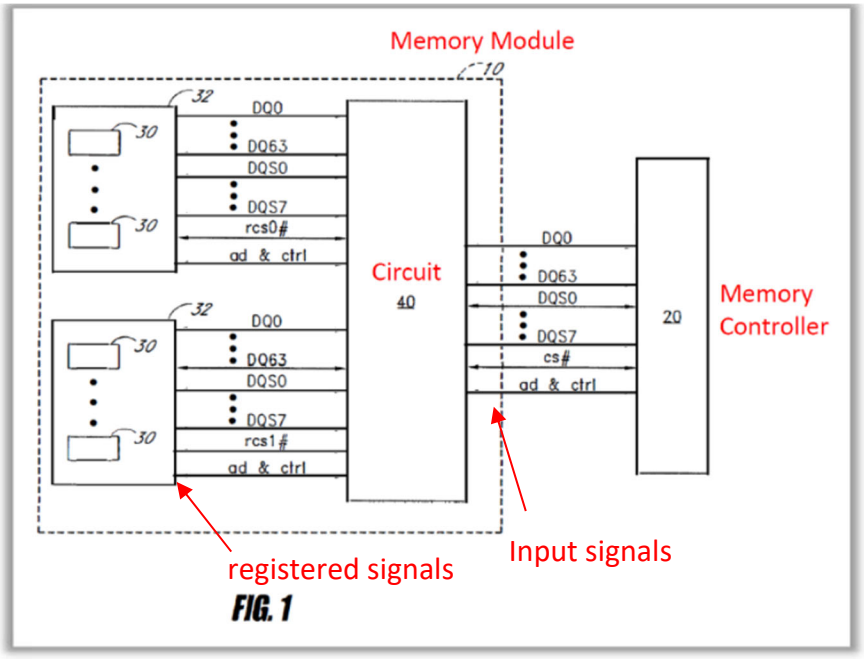

FIG. 1

75. As shown in Figure 1, "[t]he memory module 10 is connectable to a memory controller 20 of a computer system (not shown)." *Id.*, 5:67-6:1. In the Figure 1 example, the memory module 10 includes multiple memory devices 30 arranged in two ranks 32, each rank having multiple memory devices 30. *See* Fig. 1, 7:11-13. Other number of ranks are possible. *Id.*, 7:13-16.

76. The memory module 10 also includes a circuit 40 electrically coupled to the memory devices 30 and to the memory controller 20 of the computer system. *Id.*, 6:4-7. The circuit 40 includes a logic that translates between a system memory domain of the computer system and a physical memory domain of the memory module 10. *Id.*, 6:9-12. This may be achieved

by the logic receiving input data signals, chip-select signals and other control/address signals from the memory controller 20, and outputting corresponding registered data signals, registered chip-select signals and other registered control/address signals to the memory devices 30. *See* EX1001, Fig. 1, 21:66-22:35; *see also id.*, 18:37-19:9 (showing how the logic translates the input control/address and chip-select signals).

77. The memory module also includes an SPD device 240 that can report the CAS latency (CL) of the memory module to the memory controller 20. *Id.*, 22:36-38. The reported CL value of the memory module may be "one more cycle than does the actual operational CL of the memory array," for example to account for the fact that in some embodiments, "data transfers between the memory controller 20 and the memory module are registered for one additional clock cycle by the circuit 40." *Id.*, 22:38-43. The extra cycle time provides for time needed by a data buffer to perform address decoding, and/or data path multiplexing/demultiplexing, for example. *Id.*, 22:46-56.

78. The disclosure of the details of the implementation of rank multiplication is just one of the ideas found in the specification of the '417 Patent. This discussion principally is found beginning on column 22 of the specification. Ex. 1001, 22:62-34:67 (exclusive of much of the Verilog examples in columns 25-31), and this concept is not found in any of the claims of the '417

Patent. This technique—the use of two lower-capacity memory devices to simulate a larger capacity memory device for the introductory period when the newer larger capacity device is most expensive on a price per bit basis. Within the context of JEDEC-standardized memory devices and modules, this is accomplished by manipulating the addresses and chip selects so that two ranks of JEDEC-standardized SDRAMs behave like a single rank constructed of the higher-capacity devices. The values in the SPD are adjusted to hide the true organization of the memory module from the memory controller. In general, the concept is not meaningful outside the domain of JEDEC-standardized memories and their modules in which chip select signals are central to determining which DRAMs will respond to specific commands.

B. Relevant Prosecution History

79. The application that issued as the ’417 patent was filed on November 25, 2019. EX1002. There were no official rejections during the prosecution of the ’417 Patent. *Id.* The ’417 patent issued on August 17, 2021. EX1001.

C. Priority Date of the ’417 Patent

80. The ’417 Patent claims priority to five provisional applications: U.S. Provisional Patent Nos. 60/645,087 filed on January 19, 2005 (EX1005, “the ’087 Provisional”), 60/588,244 filed on July 15, 2004 (EX1006),

60/550,668 filed on March 5, 2004 (EX1007), 60/575,595 filed on May 28, 2004 (EX1008), and 60/590,038 filed on July 21, 2004 (EX1009). EX1001 at 2. The '417 Patent also claims priority to U.S. Patent Application No. 11/075,395, filed March 7, 2005, now U.S. Patent No. 7,286,436 (EX1010, "the '436 Patent"). EX1001 at 2. I understand that in connection with Ground 2, Petitioner argues that Ellsberry (filed on June 1, 2005, EX1073) is prior art to the '417 Patent because none of the five provisional applications nor the '436 Patent provide sufficient disclosure to support certain claim elements of independent claim 1. For the reasons that follow, I disagree.[2]

1. The '087 Provisional

81. The '087 Provisional relates to "memory modules for use in computers, and more specifically, relates to configurations of memory devices and other components of a memory module," and in particular to rank buffered DIMMs. Ex. 1005, p. 3. Figure 1 discloses a DIMM Memory Subsystems with two ranks and a rank buffer DIMM logic device (RBDL), which, in certain embodiments, "comprises address decoding logic which transparently translates from a system memory density domain of the computer system to a physical memory density domain of the DIMM." Ex. 1005, [0007], p. 10. This logic

[2] For purposes of this declaration, I have been asked only to render an opinion on Petitioner's arguments regarding the '087 Provisional.

translation element comprises a data path multiplexer/demultiplexer that "logically isolates the ranks of memory devices from one another," and that "advantageously diminishes the effects of jitter and other sources of noise, and simplifies various other aspects of operation of the memory module, including, but not limited to, setup and hold time, clock skew, package skew, and process, temperature, voltage, and transmission line variations." Ex. 1005, [0008].

FIGURE 1:

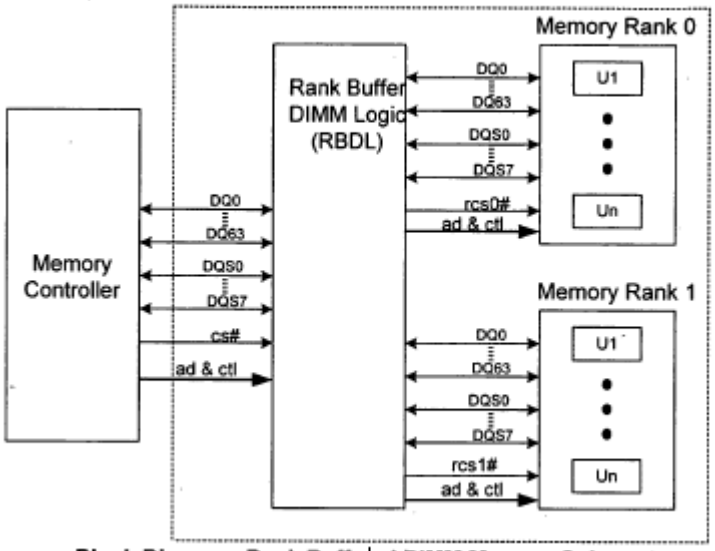

Block Diagram: Rank Buffered DIMM Memory Sub-system

Ex. 1005, p. 10.

82. The Petition's first argument is that no application filed prior to July 1, 2005 includes support for "data buffer control signals." Petition, p. 5. However, a POSITA looking at Figure 1 and paragraphs [0007] and [0008] of the '087 provisional filed in January 2005 would understand that: 1) the disclosed DIMM module has a 64-bit data bus, just as all the JEDEC-

standardized SDRAM-based DIMMs of the day; 2) the memory devices U1..Un in each rank would be burst oriented SDRAMs that together input or output a 64-bit wide burst of data in response to read or write commands; 3) the Rank Buffer DIMM Logic transfers these bursts of data between the DRAMs in the memory ranks and the memory bus connecting the module to the memory controller; 4) the data path multiplexer/demultiplexer in the RBDL would have to transfer read data bursts from one of the ranks of DRAMs to the memory bus towards the memory controller, depending on the specific read command being executed by the module (and vice versa for write commands); and 5) the claimed "circuitry" is at least in part said multiplexer/demultiplexer, which a POSITA would understand necessarily would receive control signals (i.e., the claimed "data buffer control signals") from the command decode circuitry (a POSITA would understand necessarily that the module described in the '087 Provisional would receive commands from an external memory controller and would include command decode circuitry to interpret those commands), and those control signals would be different depending on whether the read data is being transferred from rank 0 or rank 1. Only the command decode circuitry would know which rank the RBDL must direct data to or from, and so there must be some control signal from the command decode circuitry to the multiplexer/demultiplexer. Petitioner uses similar reasoning to conclude that "a

POSITA would have understood" from Perego's description of control and address information causing the selective routing of data through the buffer devices that "Perego's buffer device includes logic configured to output data buffer control signals." Petition, p. 63.

83. In light of the disclosures about the RBDL in the context of DDR2 SDRAMs and DIMMs, a POSITA would have no difficulty understanding that the disclosed multiplexer/demultiplexer in the RBDL provides sufficient support for the elements "logic … configurable to output data buffer control signals in response to the read or write memory command" and "circuitry coupled between the data signal lines in the N-bit wide memory bus and corresponding data pins of memory devices in each of the plurality of N-bit wide ranks, the circuitry being configurable to transfer the burst of N-bit wide data signals between the N-bit wide memory bus and the memory devices in the one of the plurality of N-bit wide ranks in response to the data buffer control signals." Ex. 1005, [0001], [0003]-[0010], [0013], [0017], [0019], [0022], Figure 1.

84. The Petition also says "a POSITA would have recognized that the buffer (or "circuitry") need not use "data buffer control signals" to direct a data burst to one of the ranks, even if a memory write command is directed to that one rank, because, for example, the data could be sent to both of the ranks, and a chip-select signal (e.g., rcs0#, rcs1#) could be used to target one of those

ranks." Petition, p. 8. This argument does not apply at all to reads, during which the RBDL must select the data coming from the one addressed Rank and transfer it to the memory bus. And for write commands, the RBDL must understand when to forward data from the memory bus to the rank data buses. That "when" information must be provided by data buffer control signals.

85. Petitioner also argues that a "POSITA would have also recognized that data transfer through the buffer does not necessarily happen in response to memory commands as required by the '417 Patent claims and instead could be controlled by other means, such as the 'preamble' on the DQS strobe lines." Petition, p. 8. As a preliminary matter, according to Micron's witness, Dr. Harold Stone, "some people may consider the DQS signals … also as control signals." IPR2022-00615, EX2117, 59:2-9. Under that view, which I do not agree with, a data transfer in response to DQS would also fall under the scope of data transfer in response to data buffer control signals.

86. Furthermore, I disagree that a POSITA would have viewed data transfer as being plausibly in response to DQS alone, and not in response to data buffer control signals. The DQS preamble cannot be used to control the data transfer through the RBDL. First of all, in the context of DDR2 DIMMs as disclosed by the '087 provisional there may be other modules on the memory bus. In this circumstance, it is impossible for a RBDL to differentiate a preamble

for a burst data transfer for one of its ranks from valid data being communicated to or from another rank. Even in the instance of a memory system with a single Rank Buffered DIMM, the RBDL cannot identify the beginning of the preamble for the reasons explained below. The Petition cites to Figure 23 of Ex. 1064 (reproduced below) to hypothesize the use of the preamble to control the transfer of data through the circuitry.

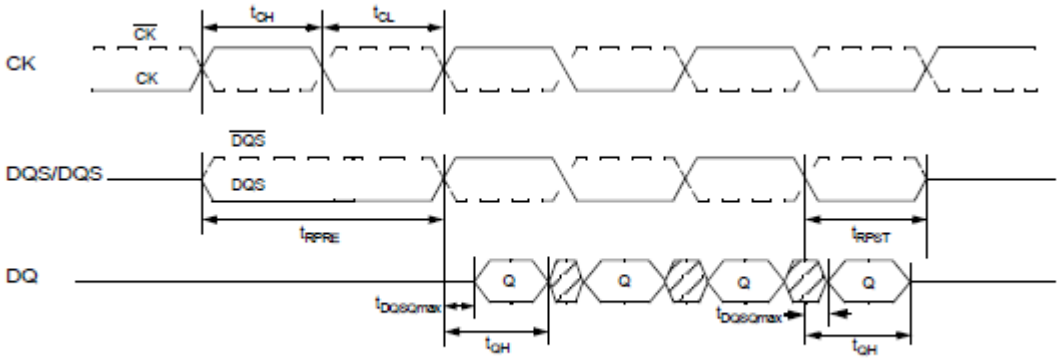

Figure 23 — Data Output (Read) Timing

87. As Dr. Wolfe acknowledged during his deposition, the DQS shown at what appears to be an intermediate level before the preamble and after the postamble does not actually indicate that the DQS signal trace is at an intermediate voltage, but only that it is not being actively driven to either a valid high voltage or a valid low voltage and thus indicates a high-impedance state with a floating voltage level. Ex. 2033 (Wolfe Deposition Transcript), 171:12-172:20. In particular, since during the post-amble the DQS signal is being driven

low, the voltage is likely to drift from low to some indeterminate voltage level between adjacent data burst transactions on the DQS signal traces.

88. Therefore, at least some of the time, there is no discernable transition in the voltage on the DQS signal at the start of the preamble and therefore nothing for the RBDL to use to identify the start of the preamble to the control of the transfer of data through the RBDL, as suggested by Petitioner. This lack of a reliably discernable preamble applies to both read and write data transfers through the RBDL. A POSITA, understanding the nature and purpose of the preambles would recognize that the data transfer through the buffer cannot be controlled in response to such preambles. Such data transfers would be controlled by data buffer control signals, which not only convey that a data transfer is to happen but also when it is to happen and what the source and destination of the transfer are to be.

89. Based on the knowledge of a POSITA, which informs the disclosures of the specification, it is clear that, contrary to the assertion of the Petitioner, the ’087 Provisional provides sufficient support for these elements.

90. Accordingly, a person of ordinary skill in the art would have understood that the ’087 Provisional discloses “logic … configurable to output data buffer control signals in response to the read or write memory command” and “circuitry coupled between the data signal lines in the N-bit wide memory

bus and corresponding data pins of memory devices in each of the plurality of N-bit wide ranks, the circuitry being configurable to transfer the burst of N-bit wide data signals between the N-bit wide memory bus and the memory devices in the one of the plurality of N-bit wide ranks in response to the data buffer control signals."

91. Thus, based on the '087 Provisional's disclosure and the knowledge of a person of ordinary skill in the art, a person of ordinary skill in the art would have understood that the inventors were in possession of the subject matter claimed by the '417 Patent. As to Petitioner's arguments regarding the lack of support for elements of Claim 1 of the '417 lack technical merit, its conclusion that the '417 Patent does not deserve a priority date before June 2005 is not supported by the evidence.

92. I understand that Petitioner has not identified any other limitations of the '417 Patent allegedly not supported by the '087 Provisional. I take that to mean that Dr. Wolfe and Petitioner concede that the '087 Provisional provides support for the other limitations. Nevertheless, for completeness, I provide below exemplary citations in the following claim chart for the remaining limitations that—when viewed through the eyes of a POSITA and the knowledge of the state of the art, including contemporaneous JEDEC standards—indicate that the inventors had possession of the claimed subject matter.

Claim/Element	Exemplary Citations of Support
1.a.1	EX1005, [0001]-[0004], [0007]-[0008], [0010]-[0011], [0014]-[0015], [0016]
1.a.2	EX1005, Fig. 1, [0001], [0004]-[0005], [0007], [0009]-[0013]-[0018], [0021]; EX1006, [0019]; EX1007 at [0015], 10, 14-16; EX1008 at [0015], 17, 21-23; EX1009, [0006]; EX1011, 4:65-66; 5:31-34, 8:58-62, 7:13-22, 33-36, 10:19-23, 10:34-37; 15:27-33, 17:17-21
1.a.3	EX1005, Fig. 1, [0010], [0018]; EX1007 at 9-11; EX1008 at 16-18
1.a.4	[1.a.1]
1.b	EX1005, Fig. 1, [0001], [0003]; EX1006, [0001], [0023], [0025]-[0026]; EX 1007, [0006], 9, 20, 22-23, 25-38; EX1008, [0001], [0007], [0019], APPX 1, [0006], 16, 27-44; EX1011, 1:22-25, 3:56-60, 7:50-54, 11:22-30, Fig. 5A-5B
1.c.1	[1.a.2], [1.a.3], [1.b] EX1005, Fig. 1, [0004]-[0011], [0014], [0016]-[0017]
1.c.2	[1.a.3] EX1005, [0005], [0008], [0009], [0017]; EX1006, Table 1, [0004]-[0005], [0009]-[0020], [0023]-[0026]; EX1007 at 10; EX1008 at 17; EX1011, 5:5-6, 6:51-55, 15:60-63, 17:31-45, 18:3-11
1.c.3	[1.c.1]-[1.c.2] EX1005, [0002], [0005], [0012], [0014]-[0016]; EX1006, Table 1, [0004]-[0005], [0009]-[0020], [0023]-[0026]; EX1007 at 10; EX1008 at 17; EX1009, [0008], [0010]-[0011]; EX1011, 5:5-6, 5:54-6:32, 6:51-55, 15:60-63, 17:31-45, 17:59-18:11
1.c.4	*See* ¶¶ 81-91, *supra*. [1.a.2], [1.a.3], [1.c.1]

	EX1005, Fig. 1, [0005], [0008], [0009], [0017]; EX1007 at 10; EX1008 at 17
1.d.1	[1.a.2], [1.b] EX1005, [0004]-[0010], [0015]
1.d.2	[1.c.2], [1.c.3], [1.d.1] EX1005, Fig. 1, [0005], [0008], [0009], [0017]
1.d.3	[1.a.2], [1.c.3], [1.d.1]-[1.d.2] EX1005, Fig. 1, [0005], [0008], [0009], [0017]; EX1007, [0015], 9-12 14-17; EX1008, [0015], 16-24; EX1008 at 16, 21, 24
1.e	*See* ¶¶ 81-91, *supra*. [1.a.2]; [1.a.3], [1.c.1], [1.d.1], [1.d.3], [1.c.4] EX1005, Fig. 1, [0005], [0008], [0009], [0017]; EX1007 at 10; EX1008 at 17
1.f	[1.e] EX1005, Fig. 1, [0005], [0008], [0009], [0017]; EX1007 at 10; EX1008 at 17
2	EX1005, Fig. 1, [0005], [0008]-[0009]
3	[1.c.4], [1.d.3], [1.e] EX1005, Fig. 1, [0017]; EX1007, [0015], 9-12 14-17; EX1008, [0015], 16-19., 21-24
4	[1.a.2], [1.d.1] EX1005, [0005], [0010], [0012], [0014]-[0016], [0019]; EX1006, [0002]; EX1007, 9-11; EX1008, [0004]-[0006], [0011]-[0014], Appx 1, 16-18; EX1011, 1:38-48, 1:57-65, 4:32-37, 4:54-59, 7:50-54, 8:1-14, 8:52-55, 12:48-13:7, Table 1, 14:4-35, 15:21-25, 15:33-43, 15:55-59, Fig. 5A-5B, Fig. 9, Fig. 10A
5	[1.c.2], [1.d.1]

	EX1005, [0002], [0005], [0012], [0014]-[0016]; EX1006, [0003]-[0020], [0023]-[0026]; EX1007 at 10; EX1008 at 17; EX1009, [0008], [0010]-[0011]; EX1011, 5:5-6, 5:54-6:32, 6:51-55, 17:31-45, 17:59-18:3-11
6	[1.e.], [1.f] EX1005, [0005], [0008], [0009], [0017]
7	[1.c.1], [1.c.4], [1.c.4], [1.e], [1.f] EX1005, Fig. 1, [0005], [0008], [0009], [0017]; EX1007, 10, 16; EX1008, 17, 23
8	[1.a.1], [1.a.2], [1.d.1], [1.d.3] EX1005, Fig. 1; EX1007, [0004], 9-10, 14-15; EX1008, [0006], Appx 1 [0004], 16-17, 21-22; EX1011, 7:29-33, 7:47-50, 10:30-34, 12:38-64, Table 1
9	[1.e.]-[1.f] EX1005, Fig. 1, [0005], [0008], [0009], [0017]; EX1006, [0009], Fig. 1; EX1007, 9-11, 17, 39; EX1008, 16-17, 24; EX1011, 7:54-56
10	[1.d.1], [1.d.3] EX1005, [0006], [0019], [0022]; EX1007, [0001], 9-10, 15-16; EX1008 at [0006], Appx. 1, [0001], 16-17, 22-23; EX1009, [0001] [00013]; EX1011, 3:56-60, 7:47-67, 10:30-37, Figs. 5A-5B
11	[1.c.4], [1.d.3], [1.e] EX1005, [0005], [0008], [0009], [0017]
12	[1.c.4], [1.d.3], [1.e.], [11] EX1005, Fig. 1, [0005], [0008]-[0009]
13	[2]
14	[8]
15	[3]

VII. OVERVIEW OF THE REFERENCES

93. I provide the following brief overview of each reference cited in the Petition.

A. Perego

94. U.S. Patent No. 7,363,422 to Perego ("Perego") was filed January 28, 2004, and issued as a patent on April 22, 2008. EX1071. Perego discloses a memory system with point-to-point topology including a "memory module having a buffer device (e.g., having a configurable width) isolating data, control, and address signals of the memory devices from the connector interface." EX1071, 4:38-45, Abstract, 5:35-55, 8:10-17, 8:20-26. Buffer device 350 is coupled to memory controller interface 375, which includes a "plurality of memory subsystem ports 378a-378n." *Id.*, 4:63-5:15. Perego's configurable width buffer device is connected to ports 378a-378n of the memory controller via a point-to-point links 320. *Id.*, 5:12-21. Perego repeatedly and consistently emphasizes the benefits of its point-to-point architecture over a conventional memory bus as found in the JEDEC DIMM standards. *See, e.g.*, Ex. 1071, 3:47-56, 4:65-5:1, 5:6-15, 5:32-55, 6:15-19, 13:49-59, 21:46-50, Figs. 3A/3B, Fig. 5A.

95. Perego's memory system (300 or 305) is built with a plurality of point-to-point buses (320) connecting a memory controller (310) with memory

subsystems (330). Ex. 1071, Figure 3A, 3B, 4:63-5:15. These memory subsystems can be inserted into sockets (380) and in these embodiments comprise modules (340, 400). Though several embodiments of buses 320 are disclosed they are all point-to-point with one end connected to the memory controller and the other to the memory subsystem. Ex. 1071, Figures 3A, 3B, 6A, 6B, 7, 8A, 8B, 8C, 4:63-5:15, 20:48-64, 20:65-21:3, 21:39-22:9. The point to point links are used to convey read and write data, and address and control information can be multiplexed with data or be transported on their own buses. Ex. 1071, Figures 5A, 5B, 5:16-31, 8:10-19, 13:49-59.

96. Perego specifically refers to a dynamic point-to-point topology to connect the memory controller to the memory subsystem/module. Ex. 1071, 3:41-47, 5:32-55, 6:57-7:29. A POSITA, being familiar with the Rambus XDR architecture, would recognize this terminology as being specific to the XDR architecture. Ex. 1071, 8:1-9, Ex. 2009, p. 10, 11, 12, 15, 17. A POSITA would recognize that the two configurations of the embodiment of Ex. 1071 at 5:32-55 describes the XDR dual module configurations in which the first module can be connected to one or two point-to-point links depending on the whether a second module is a continuity module. Ex. 2009, p. 10.

97. All of the embodiments of Perego include a buffer device and that couples the memory bus to the memory devices. Many of the embodiments

include a configurable width buffer device (391) as indicated in the title and Abstract. Ex. 1071, Abstract. The configurable width buffer device 391 is similar to buffer device 405 except that the configurable buffer device includes the Configurable Serialization units 591 and the configurable width interface 590. Ex. 1071, Figures 3C, 4A, 5A, 5B, 13:6-17, 13:49-59, 14:52-15:6, 15:31-45, 17:22-33.

98. The module/memory subsystem embodiments that are based on the configurable width buffer device 391 have a serialization ratio defined. This ratio is the ratio of the transfer rates of the memory bus and the secondary channel between the buffer device and the memory devices. Ex. 1071, 14:32-15:45. The configurable width buffer device can be programmed to adjust the width of the memory bus (the primary channel), the widths of the secondary channels to the memory devices and the serialization ratio. Ex. 1071, 14:41-51. This programmable values of the configurable width buffer circuit 391 can be set via registers in the configurable width buffer circuit 391 or control signals to the module and buffer circuit. Ex. 1071, 15:63-16:12, 16:44-17:8.

99. Perego defines his memory system in terms of "memory devices" and then goes on to list many different types of memory devices, including DRAMs, SRAMs, and DDR SDRAMs (even though these are a type of DRAM). Ex. 1071, 3:52-64. Perego specifically calls out the memory devices in the

illustrated embodiment of Figure 3C are synchronous DRAMs, which may be DDR SDRAMs, Direct RDRAMs, or XDR RDRAMs. Ex. 1071, 8:1-5. However the memory devices in the embodiment of Figure 3C may be any type of memory such as SRAMs, FRAM, MRAM, Flash or ROM. Ex. 1071, 8:5-9. These memory devices all have very different interfaces. They have different signals with different timings and voltages. A POSITA would understand therefore that the embodiment of Figure 3C is very generic and that for each type of memory device represented as a candidate for memory devices 360a, the channel 370 comprises the signal lines appropriate to that memory device type, said signal lines arranged in the topology and operating at the frequency and voltage levels appropriate for that memory type. Channel 370 as illustrated in Figures 3A, 3B and 3C do not convey any specific channel type or arrangement of memory devices. Ex. 1071, 7:65-8:9. Dr. Wolfe in his deposition acknowledged that these memory types were not directly interchangeable. Ex. 2033 (Wolfe Deposition Transcript), 131:12-23. What Perego does say about the generic channels 370 is that in the illustrated embodiment of Figure 3C, the buffer device 391 is configurable in width and that as consequence of adjustments to the memory device access width (W_A) is reduced as a result of the changing configuration of the module/memory subsystem and the configurable width buffer 391, the number of used memory signal lines may be increased or decreased. The

channels 370 are consequently subdivided into several addressable subsets. Ex. 1071, 14:52-15:6. These subsets are necessarily addressable subsets in order to satisfy the essential goal of the configurable width buffer 391: that all of the storage locations within memory devices can be accessed by the memory controller regardless of the selected memory device access width. Ex. 1071, 14:54-60.

100. The two disclosed buffer device embodiments in Figure 5A and 5B show the connections between the point-to-point link (also referred to as the primary channel (Ex. 1071, 15:7-13). Each include one or more interfaces 520 (Ex. 1071, 16:20-21). The interfaces "receive and transmit to memory devices to memory devices disposed on the module (e.g., see FIGS. 4A, 4B and 4C) via channels." Ex. 1071, 11:48-51. These are the channels 415 of Figures 4A-4C, which a POSITA would recognize as Direct RDRAM channels. Ex. 1071, Figure 4A, 4B, 4C; Ex. 2026 (Direct Rambus Tech. Disclosure), p. 8. Perego describes that the interfaces 520 can include any number of channels, as is consistent with the illustrated embodiment of Figure 4C. Ex. 1071, 11:51-55.

101. The illustrated embodiments of Figure 4A-4C utilize the buffer device 405, shown in Figure 5A. This buffer device, like the configurable width buffer device shown in Figure 5B, has a singular request and address logic block 540 which receives control and address information from the interface 510/590.

Ex. 1071, p. 13:54-59. Depending on the bandwidth of the interface 510 of Figure 5A, both interfaces 520a and 520b are combined to respond to a memory access request. Ex. 1071, 11:56-12:2. Thus interfaces are two halves of one singular memory interface that, depending on the present configuration, may be used only partially used.

102. This idea is further explained in U.S. Patent Application No. 09/797,099, entitled "Upgradeable Memory System with Reconfigurable Interconnect," filed on Feb. 28, 2001 ("the '099 Application"), which is incorporated by reference into Perego. Ex. 1071, 7:9-17. This application was issued as US patent No. 7,610,447 ("the '447 Patent"). Ex. 2034 ('447 Patent, cover page). The '099 Application discloses a module with a plurality of DRAMs 337, each of which embodies many of the features and circuits of Perego's module and configurable width buffer circuit 391. Figures 6, 7, 17, 18 and 19 of the '099 Application (i.e., the '447 Patent) show how single memory device 337 comprises a memory array 359 made up of a number array subsections 350. The memory array 359 of the memory device 337 device is connected to the memory devices I/O pins via a routing data path 354 comprised to two data path slices 360. Figures 18 and 19 and the associated text illustrate how the memory array of the DRAM is coupled to the routing data path 354, which routes data between the memory device data connections and the memory

array subsections 350. Ex. 2034 (’447 Patent), Figures 18, 19, 10:53-11:25. The modules and the corresponding memory devices have programmable width, preserving access to all of the storage locations regardless of the programmed width. Ex. 2034 (’447 Patent), 12:27-39, 12:57-65, 13:1-8.

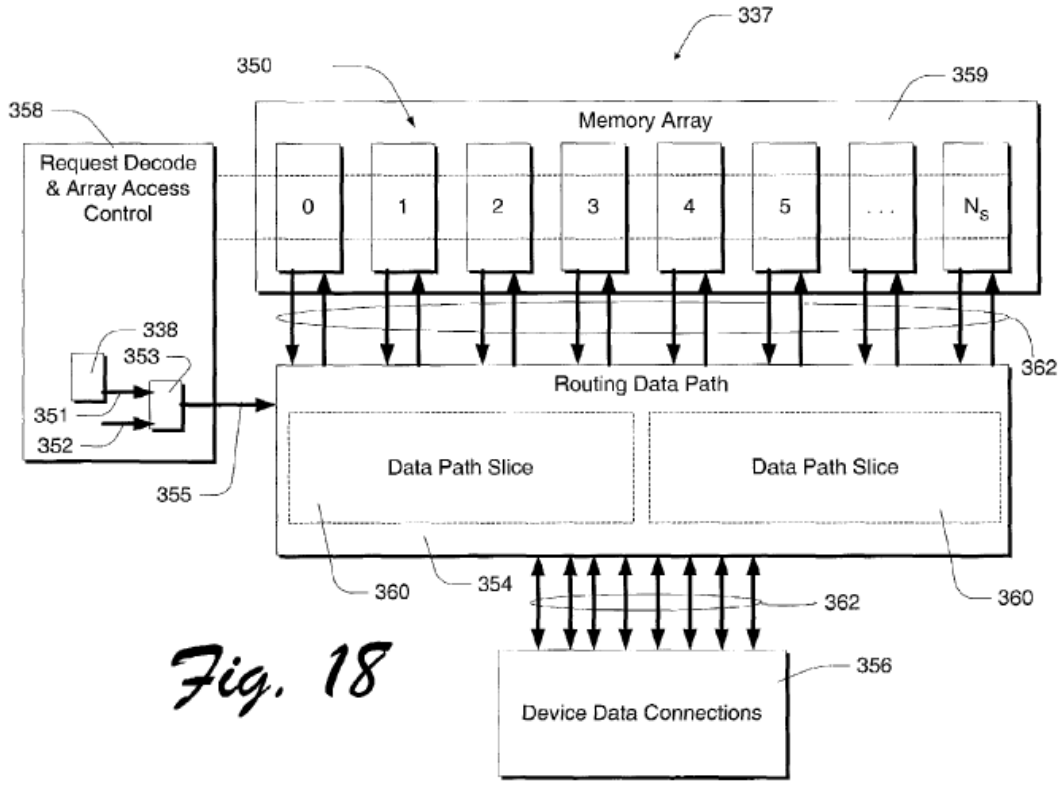

Fig. 18

Ex. 2034 (’447 Patent), Figure 18.

103. Figure 19 of the ’447 Patent shows the implementation of the routing data path of memory device 337. Ex. 2034 (’447 Patent). This circuitry allows for the width of the memory device to be programmed to be as little as one bit. Ex. 2034 (’447 Patent), 13:64-14:20. Depending on the programmed

data path width and the address of the column access command, the routing data path shunts data to or from specific memory array subsections.

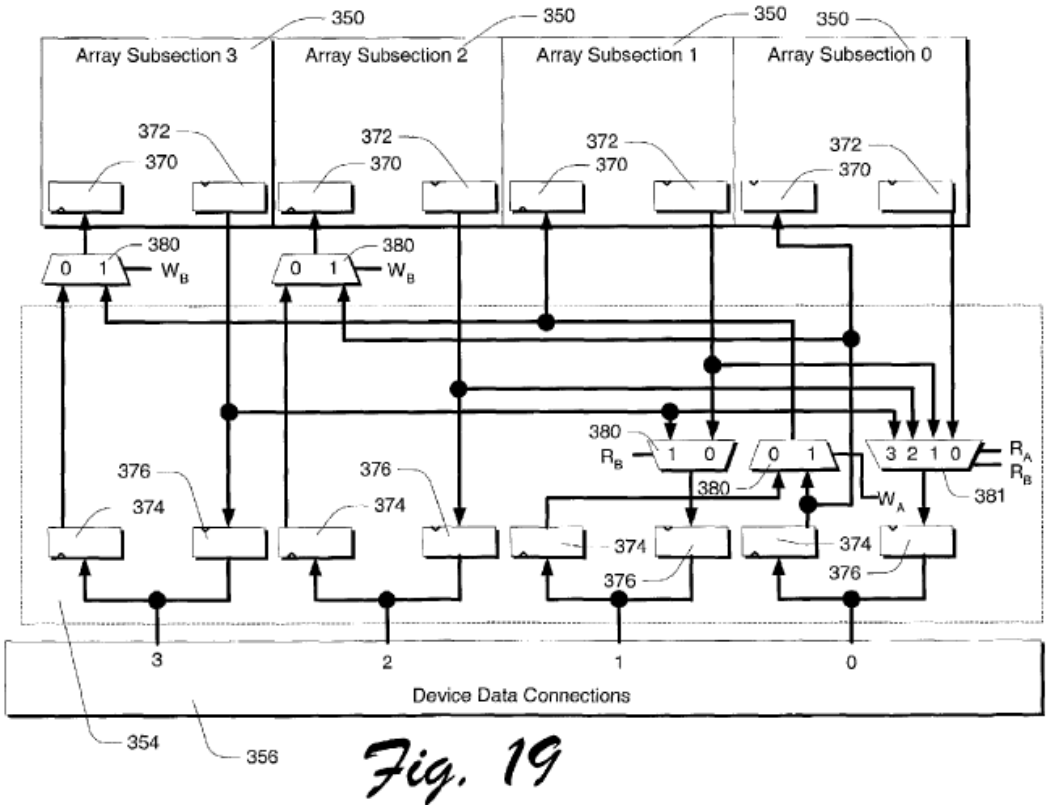

Fig. 19

Ex. 2034 ('447 Patent), Figure 19.

104. In particular, Figure 20 shows how the address bits are used in part to control the multiplexers 380 and 381 to select the read data. If the data path width is programmed to be wider than its minimum, then some of the low order address bits are ignored.

Width	Write W_A	Write W_B	Read R_A	Read R_B	Connections
1	1	1	A_0	A_1	0
2	0	1	0	A_0	0 and 1
4	0	0	0	0	3, 2, 1, and 0

Fig. 20

Ex. 2034 ('447 Patent), Figure 20.

105. These figures and description from the '099 Application incorporated by reference into Perego are very similar to the native figures in Perego: Figure 18 from the application has a lot of similarity to Figure 5B of Perego; Figure 19 from the application largely includes Figure 5C from Perego, and Figure 20 from the application is identical to Figure 5D of Perego. The two descriptions use similar terminology as well (e.g., "routing data path" vs. "configurable data path router"). Ex. 2034 ('447 Patent), 14:5; Ex. 1071, 15:33-34.

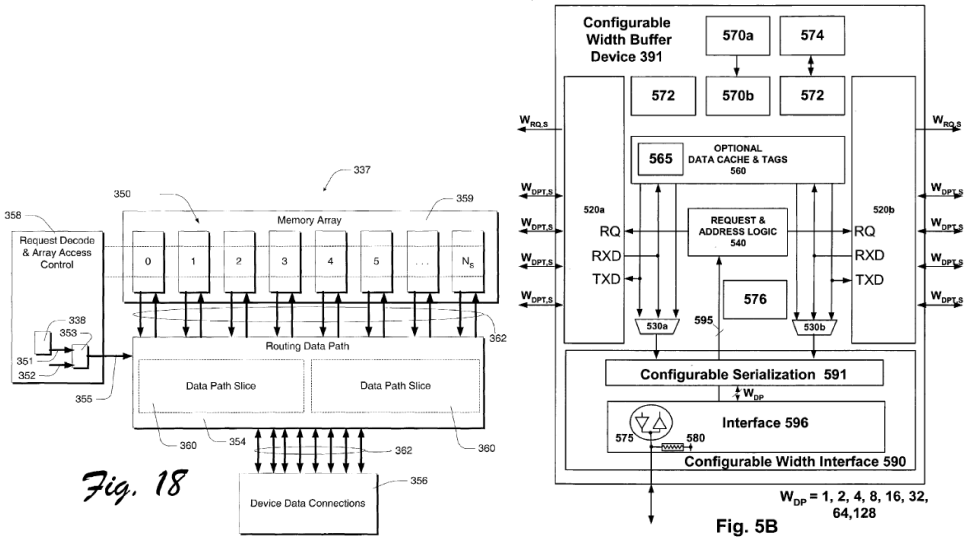

Ex. 2034 ('447 Patent), Figure 18; Ex. 1071, Figure 5B.

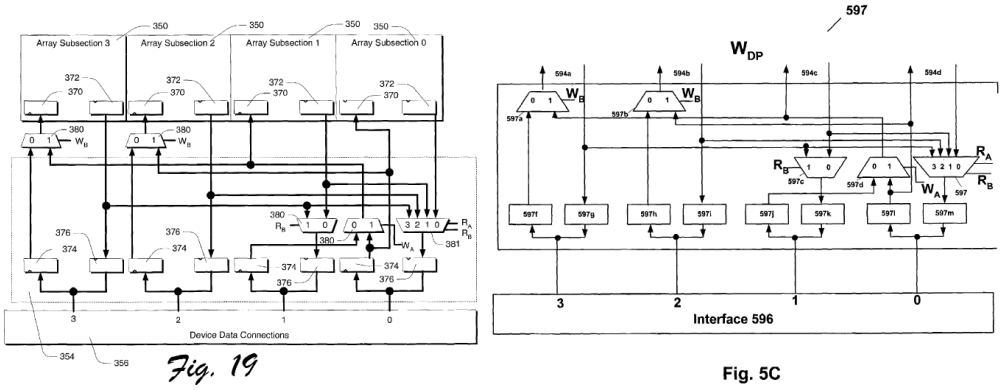

Ex. 2034 ('447 Patent), Figure 19; Ex. 1071, Figure 5C.

Width	Write W_A	Write W_B	Read R_A	Read R_B	Connections
1	1	1	A_0	A_1	0
2	0	1	0	A_0	0 and 1
4	0	0	0	0	3, 2, 1, and 0

Fig. 20

Width	Write W_A	Write W_B	Read R_A	Read R_B	Interface 596 Connections
1	1	1	A_0	A_1	0
2	0	1	0	A_0	0 and 1
4	0	0	0	0	3, 2, 1 and 0

Fig. 5D

Ex. 2034 ('447 Patent), Figure 20; Ex. 1071, Figure 5D.

106. Thus a POSITA would understand that the principals of operation of the two embodiments are similar. In particular, the '099 Application makes clear that the memory subarrays shown in Figure 18 operate as a unit and that a column access command might involve any of the subarrays, depending on the programmed device width and the address of the operation. A POSITA would understand that a prerequisite of this is that the activate command that precedes the column access must activate all of the subarrays and prepare all of them.

107. With regard then to the operation of the embodiment of Figure 5B of Perego, the native disclosure of Perego indicates that the configurable serialization unit along with the request and address logic will direct column accesses to a subset of all the secondary channels—called the target subset—according to the width of the primary link and the address of the column access, as indicated by Figure 5D. Ex. 1071, Figures 5B, 5C, 5D, 14:52-15:6, 15:31-45, 18:1-47. Perego notes that this technique of employing a configurable data path router and accessing only the target subset of secondary channels has the benefit of reducing power consumption and increasing the number of independent banks in the DRAMs on the subset secondary channels. Ex. 1071, 15:40-45.

108. As noted above, Perego discloses that memory devices 410a-410h in the illustrated embodiments of Figures 4A-4C were Rambus DRAMs. Ex. 1071, 10:54-56. These figures show channels that are linear in which the control

and data components travel in parallel to all the memory devices in the channel. This is a characteristic of Direct RDRAM channels. Ex. 1071, Figure 4A, 4B, 4C, 9:43-45; Ex. 2026 (Direct Rambus Tech. Disclosure), Figure 3; Ex. 2025 (Direct Rambus Data Sheet), pp. 46-47.

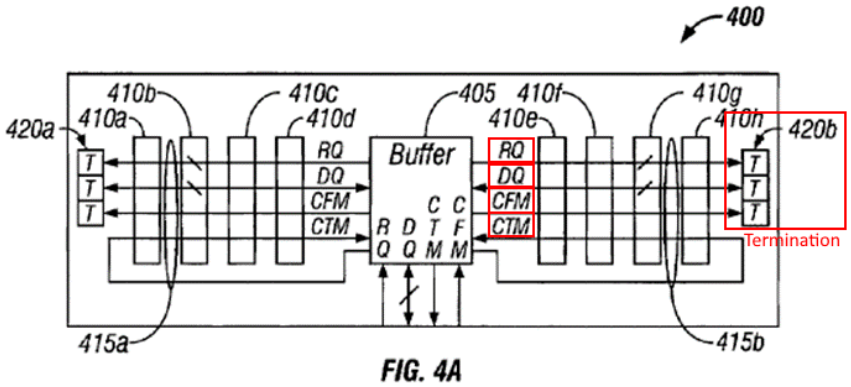

FIG. 4A

Ex. 1071, Figure 4A, Annotated.

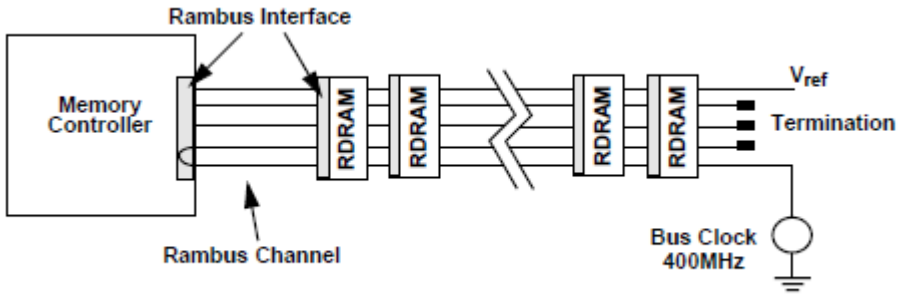

Figure 3: Primary Elements of a Rambus-Based System

Ex. 2026 (Direct Rambus Tech Disclosure), Figure 3.

IPR2023-00454
Patent No. 11,093,417

Glossary of Terms

ACT	Activate command from AV field.
activate	To access a row and place in sense amp.
adjacent	Two RDRAM banks which share sense amps (also called doubled banks).
ASYM	CCA register field for RSL V_{OL}/V_{OH}.
ATTN	Power state - ready for ROW/COL packets.
ATTNR	Power state - transmitting Q packets.
ATTNW	Power state - receiving D packets.
AV	Opcode field in ROW packets.
bank	A block of $2^{RBIT} \cdot 2^{CBIT}$ storage cells in the core of the RDRAM.
BC	Bank address field in COLC packet.
BBIT	CNFGA register field - # bank address bits.
broadcast	An operation executed by all RDRAMs.
BR	Bank address field in ROW packets.
bubble	Idle cycle(s) on RDRAM pins needed because of a resource constraint.
BYT	CNFGB register field - 8/9 bits per byte.
BX	Bank address field in COLX packet.
C	Column address field in COLC packet.
CAL	Calibrate (I_{OL}) command in XOP field.
CBIT	CNFGB register field - # column address bits.
CCA	Control register - current control A.
CCB	Control register - current control B.
CFM,CFMN	Clock pins for receiving packets.
Channel	ROW/COL/DQ pins and external wires.
CLRR	Clear reset command from SOP field.
CMD	CMOS pin for initialization/power control.
CNFGA	Control register with configuration fields.
CNFGB	Control register with configuration fields.
COL	Pins for column-access control.
COL	COLC,COLM,COLX packet on COL pins.
COLC	Column operation packet on COL pins.
COLM	Write mask packet on COL pins.
column	Rows in a bank or activated row in sense amps have 2^{CBIT} dualocts column storage.
command	A decoded bit-combination from a field.
COLX	Extended operation packet on COL pins.
controller	A logic-device which drives the ROW/COL /DQ wires for a Channel of RDRAMs.
COP	Column opcode field in COLC packet.
core	The banks and sense amps of an RDRAM.
CTM,CTMN	Clock pins for transmitting packets.
current control	Periodic operations to update the proper I_{OL} value of RSL output drivers.
D	Write data packet on DQ pins.
DBL	CNFGB register field - doubled-bank.
DC	Device address field in COLC packet.
device	An RDRAM on a Channel.
DEVID	Control register with device address that is matched against DR, DC, and DX fields.
DM	Device match for ROW packet decode.
doubled-bank	RDRAM with shared sense amp.
DQ	DQA and DQB pins.
DQA	Pins for data byte A.
DQB	Pins for data byte B.
DQS	NAPX register field - PDN/NAP exit.
DR,DR4T,DR4F	Device address field and packet framing fields in ROWA and ROWR packets.
dualoct	16 bytes - the smallest addressable datum.
DX	Device address field in COLX packet.
field	A collection of bits in a packet.
INIT	Control register with initialization fields.
initialization	Configuring a Channel of RDRAMs so they are ready to respond to transactions.
LSR	CNFGA register field - low-power self-refresh.
M	Mask opcode field (COLM/COLX packet).
MA	Field in COLM packet for masking byte A.
MB	Field in COLM packet for masking byte B.
MSK	Mask command in M field.
MVER	Control register - manufacturer ID.
NAP	Power state - needs SCK/CMD wakeup.
NAPR	Nap command in ROP field.
NAPRC	Conditional nap command in ROP field.
NAPXA	NAPX register field - NAP exit delay A.
NAPXB	NAPX register field - NAP exit delay B.
NOCOP	No-operation command in COP field.
NOROP	No-operation command in ROP field.
NOXOP	No-operation command in XOP field.

Ex. 2025 (Direct RDRAM Datasheet), p. 46, Annotated.

NSR	INIT register field- NAP self-refresh.
packet	A collection of bits carried on the Channel.
PDN	Power state - needs SCK/CMD wakeup.
PDNR	Powerdown command in ROP field.
PDNXA	Control register - PDN exit delay A.
PDNXB	Control register - PDN exit delay B.
pin efficiency	The fraction of non-idle cycles on a pin.
PRE	PREC,PRER,PREX precharge commands.
PREC	Precharge command in COP field.
precharge	Prepares sense amp and bank for activate.
PRER	Precharge command in ROP field.
PREX	Precharge command in XOP field.
PSX	INIT register field - PDN/NAP exit.
PSR	INIT register field - PDN self-refresh.
PVER	CNFGB register field - protocol version.
Q	Read data packet on DQ pins.
R	Row address field of ROWA packet.
RBIT	CNFGB register field - # row address bits.
RD/RDA	Read (/precharge) command in COP field.
read	Operation of accesssing sense amp data.
receive	Moving information from the Channel into the RDRAM (a serial stream is demuxed).
REFA	Refresh-activate command in ROP field.
REFB	Control register - next bank (self-refresh).
REFBIT	CNFGA register field - ignore bank bits (for REFA and self-refresh).
REFP	Refresh-precharge command in ROP field.
REFR	Control register - next row for REFA.
refresh	Periodic operations to restore storage cells.
retire	The automatic operation that stores write buffer into sense amp after WR command.
RLX	RLXC,RLXR,RLXX relax commands.
RLXC	Relax command in COP field.
RLXR	Relax command in ROP field.
RLXX	Relax command in XOP field.
ROP	Row-opcode field in ROWR packet.
row	2^{CBIT} dualocts of cells (bank/sense amp).
ROW	Pins for row-access control
ROW	ROWA or ROWR packets on ROW pins.
ROWA	Activate packet on ROW pins.
ROWR	Row operation packet on ROW pins.
RQ	Alternate name for ROW/COL pins.
RSL	Rambus Signaling Levels.
SAM	Sample (I_{OL}) command in XOP field.
SA	Serial address packet for control register transactions w/ SA address field.
SBC	Serial broadcast field in SRQ.
SCK	CMOS clock pin..
SD	Serial data packet for control register transactions w/ SD data field.
SDEV	Serial device address in SRQ packet.
SDEVID	INIT register field - Serial device ID.
self-refresh	Refresh mode for PDN and NAP.
sense amp	Fast storage that holds copy of bank's row.
SETF	Set fast clock command from SOP field.
SETR	Set reset command from SOP field.
SINT	Serial interval packet for control register read/write transactions.
SIO0,SIO1	CMOS serial pins for control registers.
SOP	Serial opcode field in SRQ.
SRD	Serial read opcode command from SOP.
SRP	INIT register field - Serial repeat bit.
SRQ	Serial request packet for control register read/write transactions.
STBY	Power state - ready for ROW packets.
SVER	Control register - stepping version.
SWR	Serial write opcode command from SOP.
TCAS	TCLSCAS register field - t_{CAS} core delay.
TCLS	TCLSCAS register field - t_{CLS} core delay.
TCLSCAS	Control register - t_{CAS} and t_{CLS} delays.
TCYCLE	Control register - t_{CYCLE} delay.
TDAC	Control register - t_{DAC} delay.
TEST77	Control register - for test purposes.
TEST78	Control register - for test purposes.
TRDLY	Control register - t_{RDLY} delay.
transaction	ROW,COL,DQ packets for memory access.
transmit	Moving information from the RDRAM onto the Channel (parallel word is muxed).
WR/WRA	Write (/precharge) command in COP field.
write	Operation of modifying sense amp data.
XOP	Extended opcode field in COLX packet.

Ex. 2025 (Direct RDRAM Data sheet), p. 47, Annotated.

109. This topology and signal names are not at all those of a DDR2 SDRAM on a JEDEDC-compliant memory module. Ex. 1069, Figure 19; Ex. 2001, Figure 1, pp 18-19 ("There are several classes of system nets in a

conventional memory system. For example, an SDRAM-based system may have an address net, a clock net, a data net, a DQM net, and a control net (CS, WE, RAS, CAS), (Figure 1). Each of the nets has a different loading and settling time from the other nets. A key issue limiting memory bus frequency in these systems is the fact that the loading on these nets increases nonuniformly from net to net as memory modules are added to the system.").

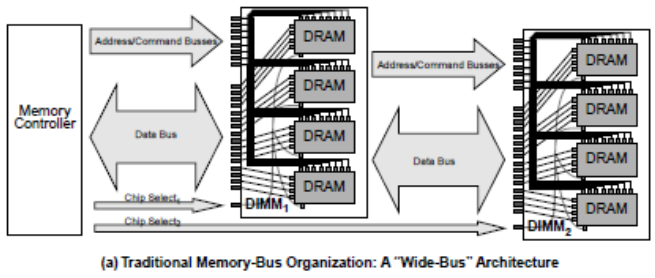

(a) Traditional Memory-Bus Organization: A "Wide-Bus" Architecture

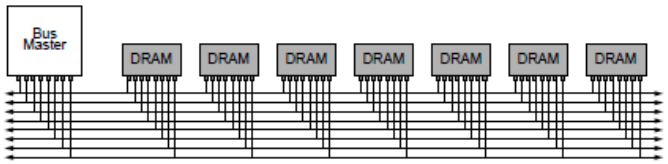

(b) Rambus Memory-Bus Organization: A "Narrow-Bus" Architecture

Figure 19: Memory-Bus Organizations
The figure compares the organizations of a traditional memory bus and a Rambus-style organization. Figure (a) shows a system of a memory controller and two memory modules, with a 16-bit data bus and an 8-bit address and command bus. Figure (b) shows the Rambus organization with a bus master and 7 DRAM slave devices.

Ex. 1069, Figure 19.

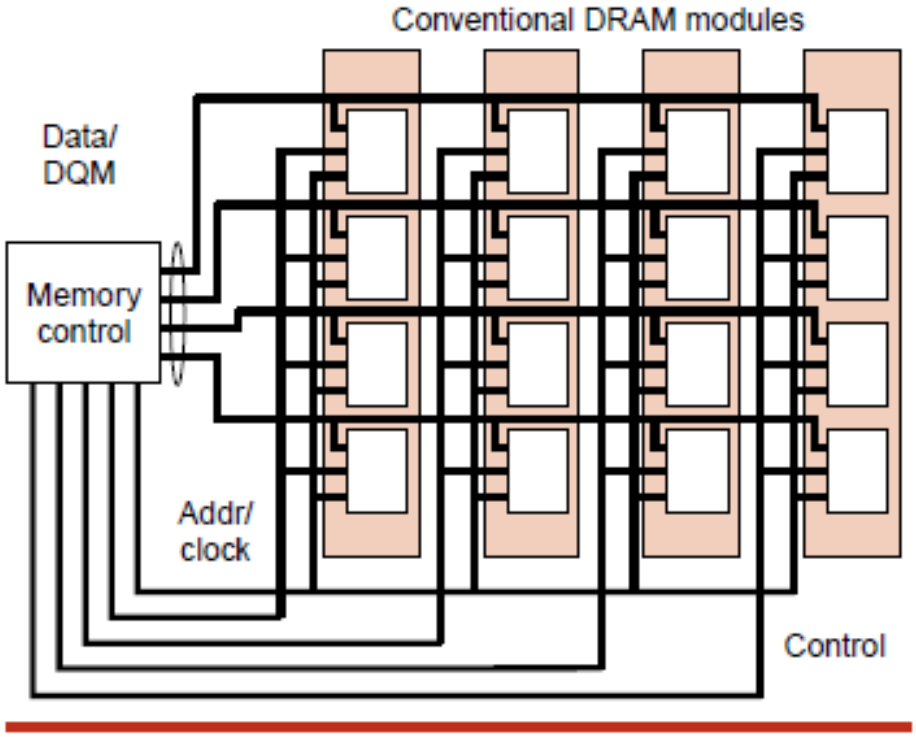

Figure 1. Nets in a conventional DRAM system.

Ex. 2001, Figure 1.

110. Thus, the topology and signal names of the secondary channels in these figures would lead a POSITA to understand that these DRAMs were Rambus Direct RDRAMs.

111. In addition, Perego also states that "other memory devices may be implemented on module 400, for example, Double Data Rate 2 (DDR2) DRAM devices and Synchronous DRAM (SDRAM) devices." Ex. 1071, 10:56-59. Perego explains that "Utilizing buffer device 405 between the memory devices and controller in accordance with the present invention (e.g., see FIG. 3) may feasibly render the type of memory device transparent to the system. Different

types of memory devices may be included on different modules within a memory system, by employing buffer device 405 to translate protocols employed by controller 310 to the protocol utilized in a particular memory device implementation." Ex. 1071, 10:59-67. In other words, Perego is not saying that devices 410a-410h would be JEDEC-compliant SDRAMs communicating directly with Direct RDRAM channels, but rather that the buffer device between the primary channel and the memory devices could be altered to communicate with DDR2 SDRAMs in a topology consistent with those memory device interfaces—i.e., consistent with Ex. 2001, Figure 1, above. Perego explains that the point-to-point primary channel would not need to be changed or downgraded to accommodate the different type of DRAMs, such as DDR2 SDRAMs, on the memory subsystem. Instead the Direct Rambus channel and the Direct DRAMs shown in Figures 4A, 4B and 4C could be replaced in another embodiment with other memory and channel types if the buffer device 405 translated the commands received on the point-to-point channel into the protocol and signals expected by the alternative memory device type. By providing a buffer device capable of this translation, the memory controller and the primary channel are independent on the specifics of the memory devices on the module, and the type of memory device on the module/memory subsystem is transparent to the system.

112. Indeed, prior art Figure 2A of Perego shows exactly this type of memory system in which the memories are SDRAM and the primary channel is a Direct RDRAM channel. "FIG. 2A illustrates a memory system **200** based on a Rambus® channel." Ex. 1071, 2:8-10.

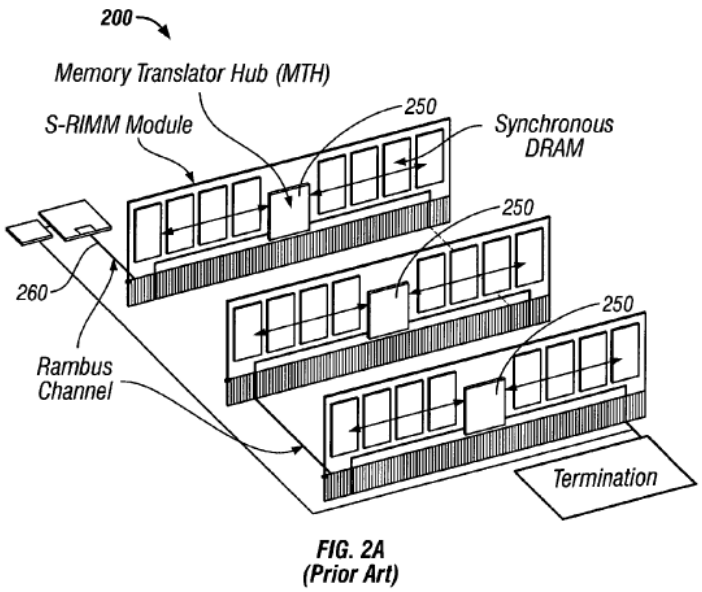

FIG. 2A
(Prior Art)

Ex. 1071, Figure 2A.

113. There is one important consequence of this independence of the memory device type that is especially relevant to the present proceeding: there are no ranks as the term is generally understood by a POSITA and no chip select signals on the primary channel. Because the memory system can include modules that include either Direct RDRAMs, which are not organized into ranks as, for example, DDR2 SDRAMs are, or include chip-select signals, as DDR2 SDRAMs it does not make sense to include this concept or these signal into the

primary channel. Indeed, since the end of each point-to-point channel 320 is a single module, and the memory locations in each module are identified purely by address bits, chip select signals are entirely superfluous. Indeed, all of the descriptions of the embodiments of the primary channel do not mention chip selects. And, for example, the modules shown in Figure 2A use SDRAM, which receive chip select signals even though the Direct Rambus Channel on which the modules sit does not include any chip select signals. Therefore, a POSITA would understand that those CS signals to the SDRAMs were generated locally by the Memory Translator Hub (MTH) from the addressing information.

B. JESD79-2

114. JEDEC standard JESD79-2, in its many versions defines JEDEC-compliant DDR2 SDRAMs. This first version of this standard was released in September 2003. Ex. 1064, Title page. These SDRAMs and the standard defining them are described above. For a SDRAM to be compliant with the standard, it must conform to all of the specifications for the physical, electrical, temporal, and behavioral characteristics of the devices in question.

115. JESD79-2 is a DRAM (i.e., memory device) standard. It is not a module standard. As such, it does not define any memory bus (or signals therein) between a module and a memory controller or any characteristics or features of any buffer or register devices on a module. Other documents ratified by the

JEDEC council relate to modules of different types and capabilities, some of which rely on the use of JEDEC-compliant memory devices.

C. Ellsberry

116. U.S. Patent Application Publication No. 2006/0277355 (Ex. 1037) ("Ellsberry") was filed on June 1, 2005 and published on December 7, 2006. Ellsberry discloses a memory module architecture that "permits transparent bank switching of memory devices." Ex. 1073, [0001]. As depicted in Figure 2, Control ASIC 204 "receives memory addresses and commands over the DIMM interface 202" from the system memory controller, and switch ASICs 206 and 208 "receive data information from the DIMM interface 202 via data buses 230 & 232, respectively." *Id.*, [0028]-[0029]. Each switch ASIC is connected to a number of memory banks (e.g., Bank 0-3 below). *Id.* Switch ASICs 206 and 208 receive respectively data byte group N and data byte group 0, provided "simultaneously" by the DIMM interface 202. *Id.*, [0030], Fig. 2 (depicting at least two data groups); Fig. 5 (depicting nine data groups).

117. Control ASIC 204 sends command signals to memory devices in different memory banks via bus 220. *Id.*, [0030]. This includes mode register set (MRS) commands that sets the CAS latency (CL) of the DRAMS in the module. See Fig. 8A for commands sent to the two ranks of memory devices.

Command	Mode	Addr	P Bank	DDR A		DDR B	
				Command	Addr	Command	Addr
MRS	X	X	X	MRS	1	MRS	1
EMRS	X	X	X	EMRS	2	EMRS	2
REFRESH	X	X	X	REFRESH	X	REFRESH	X
SELF REFRESH ENTRY	X	X	X	SLF REFRESH ENTRY	X	SLF REFRESH ENTRY	X
SELF REFRESH EXIT	X	X	X	SLF REFRESH EXIT	X	SLF REFRESH EXIT	X
SINGLE BANK PRECHARGE	Col	X	X	SB PRECHG	X	SB PRECHG	X
	Row/ Bank	X	A	SB PRECHG	X	NOP	X
			B	NOP	X	SB PRECHG	X
ALL BANK PRECHARGE	X	X	X	AB PRECHG	X	AB PRECHG	X
ACTIVATE	Col	X	X	ACTIVATE	X	ACTIVATE	X
	Row/ Bank	X	A	ACTIVATE	X	NOP	X
			B	NOP	X	ACTIVATE	X
WRITE	X	X	A	WRITE	X	NOP	X
			B	NOP	X	WRITE	X
WRITE WITH AUTO PRECHARGE	Row/ Bank	X	A	WRITEAP	X	NOP	X
			B	NOP	X	WRITEAP	X
	Col	X	X	WRITEAP	X	WRITEAP	X
READ	X	X	A	READ	X	NOP	X
			B	NOP	X	READ	X
READ WITH AUTO PRECHARGE	Row/ Bank	X	A	READAP	X	NOP	X
			B	NOP	X	READAP	X
	Col	X	X	READAP	X	READAP	X

ⓐ ⓑ ⓒ ⓓ ⓔ ⓕ ⓖ ⓗ ⓘ

Fig. 8A

Ex. 1073, Figure 8A.

118. The MRS commands received by the module from the memory controller are not sent to the ASIC switches. *See* EX1073 generally.

119. Each ASIC switch includes two data ports, Port A and Port B, coupled respectively to data buses 234 and 236. Ellsberry discloses that in column addressing scheme, memory devices coupled to both ports A and B are read but only the data from the target memory device is multiplexed onto the DIMM interface. EX1073, [0033].

D. Halbert

120. U.S. Patent No. 7,024,518 (Ex. 1078) ("Halbert") was filed on March 13, 2002, and issued on April 4, 2006. Halbert generally discloses memory module architectures and methods for operating a memory module. Ex. 1078, 1:16-19. As shown in Figures 1 and 8 in Halbert, bus 22 is arranged with a backbone of address, command, and data signal lines, which provides an electrical connection for data exchange between memory controller 20 and memory modules 100A-C. *Id*., 1:31-39, 1:61-2:14, Figs. 1, 8. Figure 8 of Halbert also shows a first rank of memory devices 140A-H on one side of the module, and a second rank of memory devices 142A-H arranged on the back side of the module. *Id*., 7:31-53.

121. Figure 7 of Halbert illustrates first m-bit wide data lines (solid lines) connected to the first rank 140A-H on the front side, and second m-bit wide data lines (dashed lines) connected to the second rank 142A-H on the back side.

Fig. 7

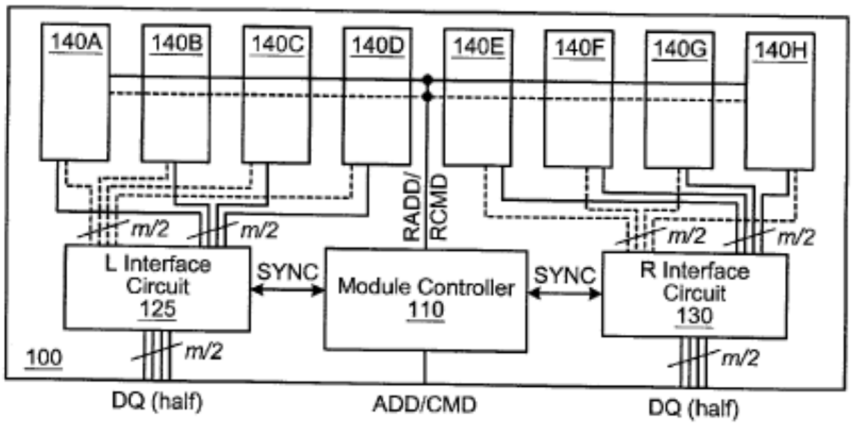

Ex. 1078, Fig. 7.

122. As shown above, the Halbert device includes module controller 110, which receives addresses and commands from a memory controller and sends those addresses and commands to memory devices 140A-H. *Id*. Halbert's memory module in Figure 7 also includes two data interface circuits (*e.g.*, left circuit 125 and right circuit 130). *Id*.

123. Halbert's Figure 4 embodiment—which Petitioner relies upon in its Petition (*see, e.g.*, Petition, p. 127)—shows the schematic representation of memory module 100 including just one interface circuit 120, and some of the internal structure of the data interface circuit 120. Data interface circuit 120 is split into two identical data interface circuits (left circuit 125 and right circuit 130) in Figure 7. *Id*., 7:31-53.

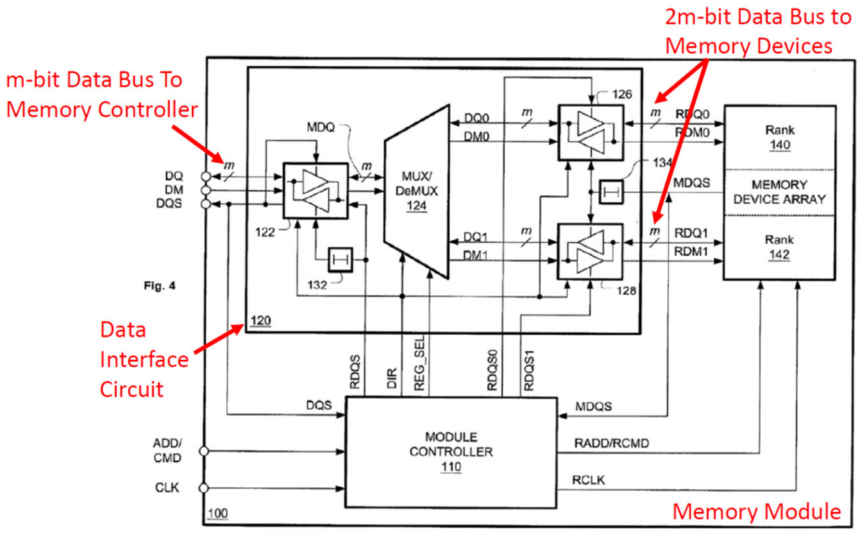

Ex. 1078, Figure 4 (annotated).

124. As shown above, memory module 100 includes module controller 110, data interface 120, and memory devices in ranks 140 and 142. Data interface 120 is coupled on the left side to the system memory controller via m-bit wide system memory data bus (DQ) and on the right side to two ranks 140 and 142 of memory devices via a 2m-bit wide module data bus (comprising two parallel m-bit wide module data buses RDQ0 coupled to rank 140 and RDQ1 coupled to rank 142). Each of the ranks 140 and 142 comprises a plurality of memory devices or a device array (e.g., rank 140 comprises memory devices denoted as 140A-140H, as shown above in Figure 7, and rank 142 comprises memory devices 142A-142H mounted on the back side of the module). Ex. 1078, Figure 8, 4:36-5:22, 7:31-53. Module controller 110 controls the operation of data interface circuit 120 by "providing timing and synchronization signals to data interface circuit 120." *Id.*, 4:40-48. Module controller 110 also provides command, address, strobe and clock signals to the memory devices in ranks 140 and 142 (*see* buses RADD/RCMD and signals RCLK).

125. Halbert teaches using the data interface circuit 120 to aggregate data transfers between "slower memory devices" (e.g., memory devices in ranks 140 and 142) and a "faster system memory bus," by accessing the multiple ranks of memory devices 140 and 142 simultaneously and in parallel. The data interface circuit 120 interfaces with "slower memory devices" via a wider (2m-

bit) module data bus at half the data rate of the narrower (m-bit) and "faster system memory bus." *Id.*, 4:36-5:22. In this manner, "***multiple ranks will receive the same address and commands***, and will perform memory operations with the interface circuit concurrently." *Id.*, 4:57-59 (emphasis added). According to Halbert, its unique design "widens the data bus on the memory module as compared to the width of the system memory data bus," which "allow[s] a faster system memory data bus to operate at full speed with slower memory devices." *Id.*, 3:43-48. In other words, the memory device array has a data interface that is twice as wide and half the data rate as the memory bus between the module and memory controller.

VIII. GROUND 1 DOES NOT RENDER THE CHALLENGED CLAIMS OBVIOUS

126. In Ground 1, Petitioner relies on a combination of Perego and JESD79-2 to find all the elements of all of the claims of the '417 Patent obvious. The analysis leading to this conclusion is flawed and misapprehends the teachings of the references, the motivation on the part of a POSITA to combine these references, and the technical characteristics of the purported combination once completed.

A. Motivation to Combine

127. Petitioner states that a "POSITA would have … been motivated to implement Perego's memory modules in a registered DIMM format with DDR

memory devices that fits into DIMM connectors and uses DIMM module input signals, according to JEDEC standards, including JESD21-C and JESD79-2." Petition, p. 32. The first flaw here is that Petitioner relies heavily on combining Perego with the portion of JESD21-C found in Ex. 1066, which is not one of the references cited as part of the combination. The RDIMM standard in Ex. 1066 relates to DDR2 SDRAMs, but the RDIMM standard is distinct from the DDR2 SDRAM standard cited as part of the combination. The RDIMM format is also functionally and operatively distinct from the Perego embodiments relied on in the Petition, as I have also discussed in detail above.

128. Furthermore, the conclusion drawn is not supported by the full disclosure of Perego in view of JESD79-2 or the portion of JESD21-C of Ex. 1066. Perego is premised on a memory system architecture and topology centered around a dynamic point-to-point interconnect connecting the memory controller to the memory subsystems/modules. Ex. 1071, 2:40-42, 4:64-5:15. The point-to-point interconnect topology allows for much faster communication and higher bandwidth solution in comparison to the multi-drop bus interconnection inherent in the design of RDIMMs. Further, Petitioner indicates that that POSITA would be motivated to transform Perego's module according to Section 4.20.10 of JESD21-C. Petition, p. 32. The JEDEC module standards specify everything needed to ensure that modules are reliably and

interchangeably available from multiple vendors, including all of the required input and output signals, timings, voltages, trace lengths and topologies. Providing a module that is built according to some, but not all, aspects of the standard, as suggested by Petitioner, is of very little value because they would not be JEDEC-complaint and thus would not be interchangeable with other JEDEC-standardized modules.

129. Regardless, Petitioner only mentions implementing "Perego's memory modules in a registered DIMM format with DDR memory devices that fits into DIMM connectors and uses DIMM module input signals, according to JEDEC standards, including JESD21-C and JESD79-2." Petition, p. 32. In this regard, Section 4.20.10 of JESD21-C requires Clock Enables, Chip Selects, exactly 3 Bank Address signals, ODT (on-die termination), bidirectional differential data strobe signals, data mask signals and a voltage reference signal. Ex. 1066, p. 6. These are all unnecessary, unusable or detrimental in the context of Perego's memory system and module architecture. For example, Perego's Direct RDRAM and XDR DRAM embodiments require additional Bank Address bits to properly address the banks internal to those memory device types, so compliance with the standard in Ex. 1066 would make the module non-functional for Direct RDRAMs and XDR DRAMs, which would defeat Perego's ability to support "[d]ifferent types of memory devices … on different modules

within a memory system," Ex. 1071, 10:63-67, as well as Perego's support of not only "backward compatibility with existing generations of memory devices," but also "new generations of memory devices," Ex. 1071, 6:39-43.

130. Perego's addressing structure is heavily influenced by the configurable data bus width of some embodiments and the reliance on Direct RDRAMs and XDR DRAMs as the DRAM types in many of the essential embodiments, and there is no suggestion in Perego that another addressing model would be consistent with the principal embodiments of Perego. In all cases, in Perego's memory system into which the disclosed modules reside, having a plurality of chip select signals on the primary channel would be superfluous because identification of the memory module and memory device is accomplished in Rambus memory structures by the address alone. As discussed above, even if DDR2 SDRAMs are used, the primary channel still follows the Rambus architecture and not the JEDEC architecture. Thus, in Perego's memory system, the buffer generates chip select signals for any DDR2 SDRAMs, and therefore providing a primary channel and modified Perego-like module that includes signal trace and socket connections for them would not make sense, and a POSITA would not be motivated to do so. This is true regardless of whether the memory devices and secondary channels are Direct RDRAM based, XDR

DRAM based, or DDR2 SDRAM based in the context of one of the protocol translation embodiments. Ex. 1071, 9:54-57, 10:54-67.

131. Petitioner argues that a POSITA would look to JESD79-2 to implement Perego's system with JEDEC-compatible memory devices. Petition p. 31. They further argue that because Perego discloses exploiting features of new generation of memory devices while retaining backward compatibility supports the idea that a POSITA would look to JESD79-2 when implementing Perego's memory subsystems. Both of these arguments are wrong.

132. First, Perego already explains how to use existing memory devices within his modules without modifying or crippling the primary link i.e. the point-to-point links 320. Perego describes how the buffer circuit can perform protocol translation in order to employ JEDEC-compliant memory devices on the modules instead of, for example, the Direct DRAM memory devices shown in Figure 4A, 4B and 4C. Ex. 1071, Figures 4A, 4B, 4C, 10:63-67.

133. Second, this technique is directly illustrated in prior art figure 2A, in which JEDEC-compliant Synchronous DRAMs are used on an S-RIMM module the interface of which is compatible with the Rambus Direct RDRAM architecture. Ex. 1071, Figure 2A, 2:8-10.

134. Third, the backward and forward compatibility is facilitated by having a high performance primary channel such as the point-to-point links

disclosed. If the module interface was downgraded to be related to the disclosure of JESD79-2, then a) it would not be possible to take full advantage of the performance of the Direct RDRAM memory devices or the XDR RDRAM memory devices and have compatibility across the full spectrum of modules contemplated by Perego (Ex. 1071, 8:1-5, 3:23-40), b) it would not be possible to use the important concept of bandwidth concentration to get the most performance possible out of Perego's inventive embodiments (Ex. 1071, 11:56-12:3) and c) forward compatibility would be lost because the primary link interface would be stuck at DDR2-levels even as DDR3 and subsequent JEDEC-standardized and Rambus DRAM architectures came to the fore.

135. Fourth, memory architecture and the primary link interface 320 shown in Perego is consistent across embodiments with buffer devices 350, 405, and 391. The configurable width buffer device in particular must adapt the address map according to the programmed primary link interface width. This is essential to preserve accessibility of all storage locations on the module in the face of dynamic changes to that width. A POSITA would soon realize that an addressing scheme based on chip selects is incompatible with such variable width interfaces because the address of storage locations would change across dynamic changes in interface width (because as the interface width changes the number of bits associated with each address changes and so the address of each

storage location changes accordingly), which is an explicit goal of Perego. Ex. 1071, 16:66-17:8. This is made even more difficult for a POSITA given the additional complication the connections from the memory to a module may also change dynamically with the change in the buffer width. Ex. 1071, 5:32-55.

136. In sum, a POSITA, considering Perego in full would realize that it is not practical to make the primary link interface resemble that of a DDR2 SDRAM as described by JESD79-2 or that of a DDR2 RDIMM as described by part 4.20.10 of JESD21-C. Instead, the XDR RDRAM interface, with its native variable width data interface and a rational addressing model is a more natural fit for the primary channel of Perego's disclosed memory module. Ex. 2009.

137. As noted above, the higher bandwidth performance of memory systems constructed of Rambus Direct RDRAMs or XDR DRAMs stems from the channel architecture, including the topology and signaling standards. Changing the primary channel of Perego to resemble that of a DDR2 SDRAM as suggested in the Petition would deprive Perego of the basis of its advantages, as well as compatibility with heterogeneous memory systems.

138. Part of the rationale provided by Petitioner for the motivation to modify the disclosed embodiments of Perego relates to the prior art module architecture disclosed in Figure 1 and elaborated on with regard to patent 5,513,135. Ex. 1071 1:58-2:6, Petition, p. 32. However, Perego is not lauding

this architecture but highlighting its deficiencies which the inventive embodiments overcome. Ex. 1071, 2:15-29, 6:24-30 ("Controller 310 receives the signals corresponding to the data at corresponding ports 378 a–378 n. In this embodiment, memory subsystems 330 a–330 n are buffered modules. By way of comparison, buffers disposed on the conventional DIMM module in U.S. Pat. No. 5,513,135 are employed to buffer or register control signals such as RAS, and CAS, etc., and address signals"). In other words, a POSITA would not read Perego as suggesting this architecture but rather as discouraging it.

139. Thus, a POSITA would not be motivated, as Petitioner suggests, to implement Perego's memory modules in a registered DIMM format … and uses DIMM module input signals according to section 4.10.20 of JESD21-C because, among other things:

1) Such a module would have lower performance than the module contemplated by Perego.
2) Such a primary channel would provide signal traces to convey signals that are not transmitted nor received.
3) Such a module would provide inadequate addressing for the memory types that are central to the inventive advantages disclosed by Perego. In particular, the module interfaces defined by section 4.10.20 of JESD21-C allow for 4 chip select signals, 3 bank address

bits, and up to 16 row and column address bits per bank. A POSITA would recognize that if the modules described by Perego were modified as described by the Petition, there would be, at a minimum, inadequate bank address bits and no interface support for sub-column addressing Ex. 2009, p. 7, 8, 11.

4) Perego already discloses the use of DDR2 SDRAMs within its modules through the use of protocol translation with certain embodiments of its buffer device. Ex. 1071, 10:63-67. Thus a POSITA would not have a reason to modify the Perego architecture away from the dynamic point-to-point interconnect system and the advantages presented by this system. Ex. 1071, 3:55-61.

5) Perego discloses that the memory interfaces 520 on buffer circuit can drive an unspecified number of subsets. Ex. 1071, 7:48-64. Perego indicates that the address of a read or write in conjunction with the current configured width of the interface, not a chip-select signal, specifies which target subset or subsets will be accessed in response to the command. Ex. 1071, 14:62-15:6. Thus an interface in which chip-select signals are an integral part of the subset selection mechanism is not consistent with the teaching of Perego.

140. The overwhelming evidence then, in light of the teaching of Perego and the limitations of the devices specified by JESD79-2 and Section 4.20.10 of JESD21-C mean that a POSITA would not alter the primary link interface to be based on these standards.

B. Independent Claim 1

1. [1.a.1] Memory Module.

141. With regard to asserting a memory module within Perego, Petitioner states that "As demonstrated by comparing Figures 3B-4C of Perego (above) with Figure 1 of the 417 Patent (below), Perego's buffered module has substantially the same structure as that of the 417 Patent: a buffer device (red) communicates with a system memory controller on a system memory bus and also with different sets of memory devices (green, blue) on the module using a respective data bus for each set of memory devices." Petition, p. 35.

142. While the coloring added to the figures of the '417 Patent and of Perego may appear to reflect superficial similarities, the presentation obscures fundamental differences between the figures and the modules the references disclose.

143. First, a POSITA would understand that Figure 1 of the '417 Patent does not disclose any physical structure. Figure 1 is schematic in nature and the two ranks will not be laid out as shown in Figure 1 but rather as shown in Figures

16A and 16B. The modules 10 of Figure 1 of the '417 Patent are meant to be compatible with standardized modules, such as RDIMMs. Ex. 1001, 6:43-50. In contrast, a POSITA would recognize that Figures 3B, 3C, 4A, 4B, and 4C of Perego show the approximate physical placement of the devices on the module, with the buffer devices shown in the middle of the module as was common in such modules so that signal trace lengths can be minimized and equalized. Put simply, a central buffer does not imply the existence or number of ranks or anything regarding the rank structure.

144. Second, while the '417 Patent seeks to preserve commonality with the JEDEC standardized modules, Perego explicitly turns away from that style of memory system architecture to develop systems that are both higher performance and expandable and which different embodiments can use different memory types without changing the primary channel characteristics or the advantages of its memory system architecture. Ex. 1001, 6:43-50, Ex. 1071, 4:63-5:15.

2. [1.a.2] Operable in a Computer System

145. On page 36, Petitioner misrepresents the motivation behind Perego's disclosed architecture and embodiments. Specifically, Perego's Figures 1, 2A and 2B represent prior art embodiments that Perego's embodiments allegedly improve upon. The memory system architecture of Figure 1 is a

conventional memory system using modules that, like the JEDEC-standardized modules, have shared multi-drop control, address and data buses (130 and 140) and separate module control lines 150 that are dedicated to each module in the memory system. Ex. 1071, 1:31-41.

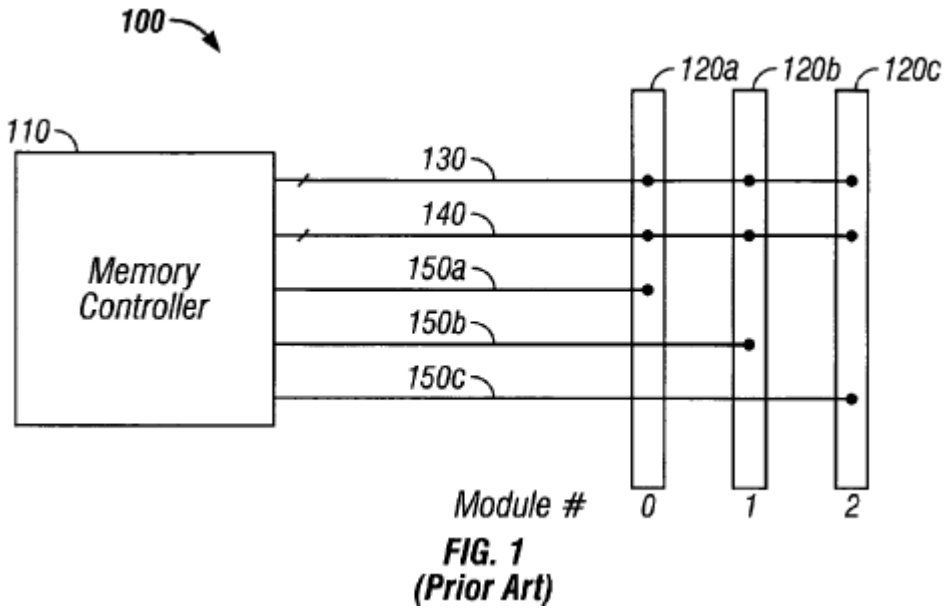

FIG. 1
(Prior Art)

Ex. 1071, Figure 1.

146. This prior art architecture provides an upgradeable memory system and configurations with different memory capacities. Ex. 1071, 2:15-17. But each configuration may present different electrical characteristics to the control/address bus. Ex. 1071, 2:17-22. A POSITA would understand that this variation in the electrical characteristics limits the performance of the memory system, Ex. 1071, 2:22-25.

147. The need that Perego's disclosed architecture and embodiments address is for a memory system architecture and implementations that replicate the upgradeability of the prior art system of Figure 1 but at higher bandwidth, principally by using point-to-point interconnects in which the electrical characteristics are constant regardless of the configuration of the memory system. Ex. 1071, 2:26-29. This does not in any way, imply compatibility with prior-art computers. Indeed since Perego's memory architecture, as shown in Figures 3A and 3B are very different from that in the prior art Figure 1, there is no suggestion of backward compatibility with that class of system.

148. The Petition further misrepresents the teaching of Perego with regard to this element on page 38. The paragraph there, referring to the embodiments that use configurable width buffer 391 shown in Figure 5B, states that in these embodiments, the disclosed data bus is 128 bits. That is the number of data pins on the buffer circuit and the edge of the memory module. W_{DP} is the configured or programmed subset of the data bus that is used at a particular time. The described implementation of the configurable width buffer device is that the portion of the data bus that is used can be changed dynamically. In these embodiments, the width of the data bus is 128 bits, not all of which are necessarily used, depending on the dynamic configuration of the buffer device. Ex. 1071, 14:16-31.

149. In both the embodiments using the buffer device 405 (shown in Figure 5A) or the configurable width buffer device 391 (shown in Figure 5B), the multiplexers 530a and 530b perform "bandwidth-concentrating operations" to account for different widths and transfer rates on the memory device side (interfaces 520a and 520b combined) and the memory bus side of the buffer device (interface 510 in Figure 5A and interface 590 in Figure 5B). Ex. 1071, 11:56-12:3, 15:7-45, 16:17-21, 21:4-38.

3. [1.a.3] Signal Lines

150. Claim 1 of the '417 Patent requires that the memory bus connecting the module to the memory controller comprise address and control signal lines and data signal lines. The petition cites to point-to-point links 320 for this element. But point-to-point link 320 is described as being a multiplexed bus in which address, control and data are all sent sequentially over the same signal lines, requiring these different pieces of information to be separated within the buffer circuit. Ex. 1071, 13:49-59. Instead, the Petition points to passages in column 9 of Perego which are referring not to the memory bus but to the connections 415 between the buffer circuit and the memory devices. These passages are alluding to the embodiments in which the memory devices can be DDR2 SDRAMs and the buffer circuit is providing the necessary protocol translation to facilitate the different interface at the memory bus and memory

device levels. Ex. 1071, 9:34-64, 10:63-67. Petitioner also notes that Figures 4A and 4B show RQ and QD lines connecting to the memory bus, but fails to note that in these embodiments, the primary channel includes not only RQ and DQ but clock signals CTM and CFM. Petition, p. 41 (citing Ex. 1071, 9:58-60, Figs. 4A, 4B). A POSITA would understand, as described above, that these signals are those of a Direct RDRAM channel. This is what is shown in the prior-art embodiment of Figure 2A where a single bus-structure Rambus Channel is used to connect high-bandwidth channels as described in the Direct RDRAM architecture. Ex. 2001, pp. 26-27. These bus embodiments for the primary channel are different that those relied upon by Petitioner for the memory bus element [1.a.2]. See Petition, pp. 36 (asserting that Perego's module "is compatible with prior-art computers using a traditional bus for data, address, and control signals consistent with the JEDEC standards"), 38 ("A POSITA would have understood that a data width of W_{DP}=64 corresponds to the 64-bit data width of a JEDEC-compliant registered DIMM module.").

4. [1.c.1], [1.c.2] and [1.c.3] The Set of Received Address and Control Signals Include Input Chip Select Signals and the Registered Address, Control and Chip Select Signals

Correspond to the Received Address, Control and Chip Select Signals.

151. With regard to the analysis of elements [1.c.1] (input address and control signals to the logic) and [1.c.2] (input chip select signals to the logic), the Petition makes similar errors as above.

152. Claim 1 of the ’417 patent requires that the claimed logic on the module receive address and control signals, the control signals including a plurality of chip select signals, and that the same logic outputs corresponding registered copies of those signals.

153. The Petition’s analysis of these elements relies heavily on the assumption that a POSITA would modify Perego’s primary link interface to be at least similar to a DDR2 memory device interface, as defined by JESD79-2 and consequently this modified version of Perego’s module would receive address and control signals, and that those control signals would include a plurality of input chip select signals.

154. As described above, this assumption is invalid. On careful consideration of the teaching of Perego, a POSITA would not adjust the primary link (the memory bus interface) to mimic that or a DDR2 SDRAM or that of a DDR2 RDIMM. Any attempt to do this by a POSITA would change the operating principles of Perego regarding how to provide a high performance memory module that provides backward and forward compatibility and

consistent addressing of storage location in the face of variable and dynamically configurable data path width and the memory device independence articulated by Perego. Ex. 1071, 10:59-63 ("Utilizing buffer device 405 between the memory devices and controller in accordance with the present invention (e.g., see Fig. 3) may feasibly render the type of memory device transparent to the system.").

155. As explained above, a POSITA would not modify the teaching of Perego to change characteristics of the primary link interface to be similar to that of a DDR2 SDRAM or a DDR2 RDIMM. Instead, Perego teaches that the addressing model of Perego's module is flat: there are no chip select inputs to the module in any of the embodiments and, as explained in more detail below with regard to element [1.e] the memory devices populate a flat address space. The disclosure of the '099 application incorporated by reference and the resulting '447 Patent confirm this.

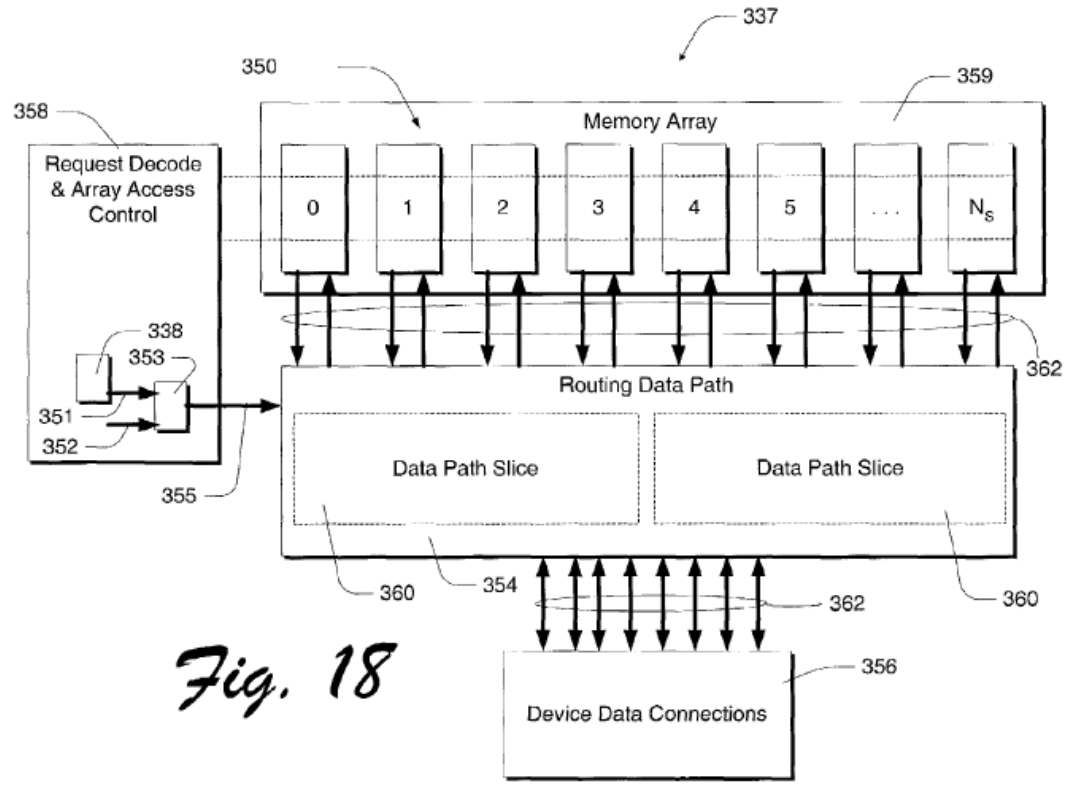

Fig. 18

Ex. 2034 ('447 Patent), Figure 18.

156. The teaching of Perego, supported by that of the '447 Patent is that the two interfaces 520a and 520b operate as a unit controlled together by the central Request & Address Logic block 540. Ex. 1071, 13:49-59. A POSITA would understand that the interface 520a is for memory devices on left side of the buffer device (viewed from the front) and interface 520b is the interface for memory devices on the right side of the buffer device on the module printed circuit board. A POSITA would recognize this physical and logical orientation because it was the norm at the time, including, for example, among all the RDIMMs specified in Section 4.10.20 of JESD21-C. Ex. 1066, pp. 44-48.

157. Each interface 590 has one or more ports for controlling one or more channels for a total of two or more channels per buffer device. Ex. 1071, 11:48-55. As discussed above, a POSITA would understand that the nature and topology of the channel depends on the type of memory being controlled (e.g., DDR2 SDRAM vs Rambus Direct RDRAM). In some embodiments, the pins of the channels are configurable so that one buffer device design can control a variety of different DRAM types. Ex. 1071, 13:60-14:15. Finally, depending on the address of a read or write, the size of the data transfer, the buffer device data path width and the serialization ratio, a read or write command is sent to one or more ports. Ex. 1071, Figures 5C, 5D, 15:37-40, 18:1-47.

158. A POSITA would understand that because the targeted subset for a read or write command is selected based on the column address, all of the DRAMs in all of the target subsets have to be ready to receive a read or write command. A POSITA, knowing how DRAMs are controlled, would understand that all of the DRAM would have to be activated identically so that each of them would be receptive to a read or write. In this sense the corresponding DRAMs on each of the targeted subsets are controlled as a single unit and operate Activate commands together. Ex. 1064, pp. 9-10, 23-41.

159. To the extent that a Perego module populated with DDR2 SDRAMs is a separate embodiment distinct from the others and potentially with a unique

interface design, Perego does not suggest such an interpretation. Specifically, Perego teaches that modules with different memory types, each implementing the appropriate protocol translation are expected to cohabitate in a single memory system. This would not be possible if the SDRAM populated module has a different interface design than, so example, the Direct RDRAM or XDR RDRAM populated modules. Ex. 1071, 10:63-67.

160. Petitioner argues that Perego teaches rank multiplication and that the language about target subsets being selected by addresses teaches the pervasive use of rank multiplication and the increase in presumed chip select signals at the memory device side of the buffer circuit relative to the primary link side. Petition, pp. 50-51. This not correct. Perego does not implement rank multiplication. There is no suggestion within the disclosure of Perego that the memory controller is presented with a view of the memory system in which it appears to the memory controller that two or more memory devices are actually one larger device. The idea that addresses can be used to identify and route transactions to specific target subsets of channels 370 is completely different from the rank multiplication that can occur in memory modules of the style standardized by JEDEC such as those found in JESD21-C. Ex. 1071, 11:56-12:3, 15:31-42; Ex. 1066.

161. This idea of addressable target subsets of a wider channel in the embodiments with a configurable width buffer device 391 stems from the fact that in these embodiments the configured width of the active portion of the primary link 320 can be less than the width of the channel 370 and address bits are how the buffer device 391 knows which portion of the wide channel 370 is being accessed. Ex. 1071, 7:30-64. Data is routed to and from the narrower portion of the primary link 320 to the addressed portion of the wide channel 370 by the configurable data path router. Ex. 1071, 15:31-42. This idea of narrow subsets of a memory array is evident in the sub-column address bits found only in XDR RDRAMs; for the other memory device types, the width of an addressable target subchannel is limited by the width of the memory devices. Ex. 2009 (XDR Arch.), pp. 2, 6, 8, 12, 16. In JEDEC-standardized modules, the individual byte lanes are not allocated unique addresses that the module is cognizant of. Ex. 1066, pp. 6-7, 42.

5. [1.d.1] Ranks

162. Claim 1 requires that the recited memory devices are mounted on the printed circuit boards and arranged in a plurality of N-bit wide ranks. Ex. 1001, 42:41-42. Further, the claim requires that the plurality of N-bit wide ranks correspond to the plurality of registered chip select signals such that each of the plurality of registered chip select signals is received by one of respective N-bit

wide ranks. Ex. 1001, 42:42-47. Finally, the claim requires that one of the N-bit wide ranks is configured to receive or output a burst of N-bit wide data signals in response to the read or write command. Ex. 1001, 42: 48-53.

163. Petitioner's proposed construction of rank is "an independent set of one or more memory devices on a memory module that act together in response to command signals, including chip-select signals, to read or write the full bit-width of the memory module." Petition, p. 26.

164. As noted above, a district court in a related proceeding has construed a rank to be "a 'bank' of one or more devices on a memory module that operate in response to a given signal." Ex. 2030 (Claim Construction Order), p. 15. As explained above, this construction is not consistent with memory modules in existence around the time of the '417 Patent or Perego (or subsequently to this day) due to the ambiguous term "bank" and—if "a signal" means a single signal—the requirement that memory devices operate in response to a single signal. It is not consistent with Petitioner's proposed construction, which expressly requires chip-select signals and reading or writing the full bit-width of the memory module (Petition, p. 26), for example, and also conflicts with Dr. Wolfe's description of "rank" as including memory devices that "are acting together in response to the chip select signals, to read or write the full bit

width of the memory module." Ex. 2033 (Wolfe Deposition Transcript), 154:7-24.

165. Under Petitioner's proposed construction, Petitioner fails to demonstrate that Perego discloses the elements of claim 1 of the '417 patent relating to a plurality of N-bit wide ranks, and in fact Perego does not disclose these elements. Petitioner begins by quoting Perego as disclosing "grouping memory devices into multiple independent target subsets (i.e. more independent banks)," Petition, p. 67 (quoting Ex. 1071, 15:37-45), as "render[ing] obvious the claim limitation 'ranks' as properly construed." Petition, p. 67. Dr. Wolfe likewise testified in this proceeding that the same passage of Perego references ranks. Ex. 2033 (Wolfe Deposition Transcript), 71:5-72:4. But this passage has nothing to do with ranks, much less ranks according to Petitioner's proposed construction. Indeed, in a prior deposition regarding a related Netlist patent, Dr. Wolfe testified: "I do not think that this paragraph [the same paragraph in Perego] specifies how ranks are selected or enabled," and he further testified that "independent banks" refers to "the number of page registers which can be activated corresponding to different addresses" "throughout the system." Ex. 2003, 202:8-204:10. Indeed, this paragraph of Perego states that "the target subset of secondary channel lines may be selected via address bits," which are distinct from the chip-select signals required by Petitioner's own proposed

construction. See also Ex. 1069, p. 2 (stating that "the chip-select bus, is essential in a JEDEC-style memory system").

166. In the analysis of the rank elements, Petitioner also looks to the illustrated embodiments of Figures 3C, 4A, 4B, and 4C and colors them to suggest that these embodiments disclose ranks generally. I disagree.

167. Looking first to the illustrated embodiments of figures 4A and 4B. As explained above, a POSITA would understand that these embodiments are Direct RDRAM embodiments and that the illustrated memories 410 of Figure 4A and the unlabeled memory devices of Figure 4B are DRDRAMs and not DDR2 SDRAMs or any other variety of JEDEC-standardized DRAM memory devices. Ex. 1071, Figures 4A, 4B, 2:46-47, 9:26-57, 10:5-13. Direct RDRAMs do not receive a chip-select signal. Instead, as explained above, Direct RDRAMs receive commands that include a 5 bit chip ID as integral part of the packetized commands. Direct RDRAMs are programmed with a unique ID and thereafter they recognize and act on commands that include that ID. Ex. 2001, p. 20, Ex. 2026 (Direct Rambus Tech. Disclosure) p. 11; Ex. 2025 (Direct RDRAM Data Sheet), pp. 2, 6-9, 29.

168. In addition, the illustrated Rambus Direct RDRAM data channels are 16 or 18 bits wide only. Ex. 2001, pp. 20-21. Petitioner's argument is based

on finding a 64-bit wide data bus in the primary link and the same width at the memory devices. This is certainly not the case here.

169. Therefore for at least these reasons, the illustrated embodiments of Figures 4A and 4B of Perego do not disclose the rank elements under Petitioner's proposed construction of rank.

170. A POSITA in considering the illustrated embodiment of Figure 4C juxtaposed with those of Figure 4A and 4B and using the same unit numbers for the channels and the memory devices, along with the topological similarities to Figure 4B, would indicate to a POSITA that this embodiment also uses Direct RDRAMs. Ex. 1071, Figure 4B, 4C, 4A, 4B, 2:46-47, 9:26-57.

171. Petitioner also looks to the illustrated embodiment of Figure 3C to find the claimed ranks within Perego. Figure 3C is a more generic figure that does not indicate the type of DRAMs being used within it. Perego does indicate that the configurable width buffer device 391 shown in Figure 3C and 5B can perform protocol translation to facilitate memory modules using legacy memory devices such that the modules are compatible with the point-to-point primary link without modification. Ex. 1071, 2:43-45, 7:30-39, 7:65-8:5, 13:6-17.

172. As described above, a POSITA considering the embodiment of Figure 3C would understand that the interfaces 520a and 520b of configurable width buffer device 391 are part of a pair working in tandem, controlled by the

Request and Address Logic 540. In a further embodiment, the pair of interfaces 520a and 520b can together can be programmed to access the JEDEC-standardized memory devices having widths between x4 and x16 totaling 64 bits. Also as explained above, A POSITA would recognize this physical and logical organization as sharing this feature with the JEDEC-standardized RDIMMs, which include a register device in the middle of the module. This understanding is reinforced by the disclosure of the '099 Application incorporated by reference and the use of addresses to select target subsets for read and write transactions. Also as described above, this means that all of the memory devices on the target subsets, recognized by a POSITA in Figure 5B as the up-to-8 channels of width $W_{DPT,S}$ need to be activated together, meaning that they are all part of one DDR SDRAM rank.

173. When the disclosure of Perego is considered in its entirety, a POSITA would conclude that the illustrated embodiment of Figure 3C discloses, in a DDR2 SDRAM configuration, a single rank of up to 64 bits in width. However, the configurable width interface 590 of configurable width buffer device 391 illustrated in Figure 5B has width 128 bits regardless of its being configured at any one time to use any binary number of these from 1 to 128. Ex. 1071, 14:19-23.

174. Therefore, Petitioner's finding of the rank element of Claim 1 of the '417 Patent fails because 1) Perego, in its disclosed Direct RDRAM embodiments fails to disclose the required chip selects or the same bus width at the memory devices and the primary channel, and 2) in its DDR2 SDRAM embodiment of Figure 3C, Perego fails to disclose a plurality of ranks each with its own registered chip select signal that corresponds to an input chip select signal on the primary channel and it fails to disclose the same bus width at the memory devices and the primary channel.

175. If the Board finds that Figures 4A, 4B, and 4C show memory devices 410 that can be DDR2 SDRAMs as Petitioner asserts, then the rank elements of claim 1 of the '417 patent are still not present in Perego. With regard to Figure 4A, Petitioner asserts that both the collection of memory devices 410a through 410d and the collection of memory devices 410e through 410h disclose a 64-bit wide rank. But as described above, a POSITA would understand from the full disclosure of Perego that the two interface units 520a and 520b operate together and so if indeed each of these collections offers a 64-bit data path to the buffer device, the structure that is disclosed is a 128-bit rank not two 64-bit ranks. But Petitioner argues that it would have been obvious for a POSITA to modify the modules disclosed by Perego to have a JEDEC-compatible 64-bit data bus. Thus, even in this circumstance, the ranks of Perego Figures 4A, 4B

and 4C do not match the width of the module's memory bus as required by Petitioner's own proposed construction of "rank" in claim 1.

6. [1.e] Circuitry Coupled Between Data Signal Lines and Corresponding Data Pins

176. Element [1.e] of claim 1 of the ''417 patent recites in full:

> circuitry coupled between the data signal lines in the N-bit wide memory bus and corresponding data pins of memory devices in each of the plurality of N-bit wide ranks,
>
> the circuitry being configurable to transfer the burst of N-bit wide data signals between the N-bit wide memory bus and the memory devices in the one of the plurality of N-bit wide ranks in response to the data buffer control signals and
>
> in accordance with an overall CAS latency of the memory module;

177. As explained above, Perego fails to disclose a plurality of N-bit wide ranks under the Petitioner's own proposed construction. Therefore, this element too is not disclosed by Perego under that construction of "rank".

7. [1.f] CAS Latency

178. Element [1.f] recites in full "wherein data transfers through the circuitry are registered for an amount of time delay such that the overall CAS latency of the memory module is greater than an actual operational CAS latency of each of the memory devices."

179. This element explicitly requires data transfers to be registered through the claimed circuitry en route between the memory devices and the

primary link. Thus, the claimed time delay must include delay in the data path. Delay on the address/control path alone cannot satisfy this element. The JEDEC-compliant RDIMMs defined in JESD21-C include no "circuitry" between the memory devices and the module inputs by which data may be registered. See Ex. 1066 (JESD21-C), p. 9. Thus, when Petitioner talks at length about RDIMMs and the delay due to registering the address and command signals, citing to Ex. 1062 (JESD21-C), p. 68, it is irrelevant to this claim element. Petition, p. 95. Furthermore, this disclosure is found only in the DDR RDIMM specification of JESD21-C, which is not asserted by Petitioner in any ground, and which a POSITA would not be motivated to implement in connection with Perego for the reasons explained above.

180. Similarly, references to Additive Latency, are equally irrelevant. The Petition says the purpose of Additive Latency is to "leave enough time for the [address and control] register on the memory module to perform its functions." Petition, p. 95. This is completely false. Additive Latency relates to the latency of the SDRAMs (so it is part of the "actual operational CAS latency of each of the memory devices"). It is not about any additional latency added at the module level. It was added to the definition of JEDEC-standardized SDRAMs in order to allow row commands and column commands to be delivered back to back, thereby greatly simplifying the scheduling of the

memory device command and address bus. Ex. 1064 (JESD 79-2), p. 24 ("[T]he DDR2 SDRAM allows a CAS read or write command to be issued immediately after the RAS bank activate command (or any time during the RAS-CAS-delay time, tRCD, period). The command is held for the time of the Additive Latency (AL) before it is issued inside the device."). This is why the AL parameter is programmed into the SDRAM device itself using an EMRS command and not on a module component. Ex. 1064 (JESD79-2), p. 14.

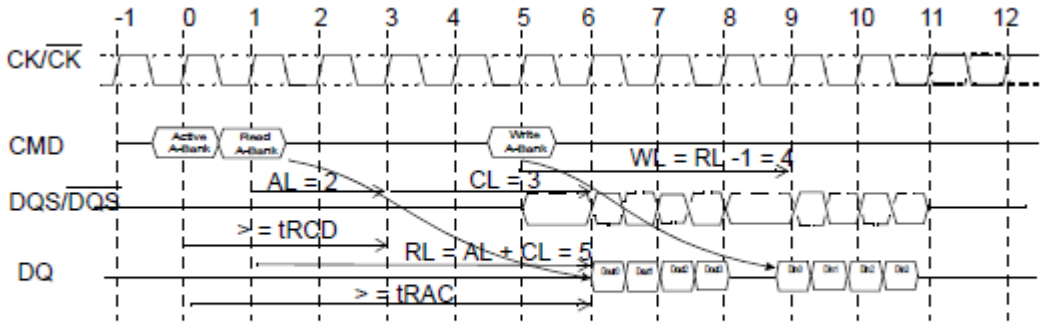

Figure 21 — Example 1: Read Followed by a Write to the Same Bank [AL = 2 and CL = 3, RL = (AL + CL) = 5, WL = (RL - 1) = 4, BL = 4]

Ex. 1064 (JESD79-2), p. 24.

181. The Petition goes on to say that "It would have been obvious to add one additional clock cycle to Perego for the same reason, i.e., so that the memory module complies with the timing of the JEDEC standards." This is equally false and completely misrepresents the incompatibilities between the JEDEC-standardized memory modules and the modules contemplated by Perego. For example, the Direct RDRAMs depicted in Ex. 2026 (Direct Rambus Tech.

Disclosure) have a read latency of 8 cycles (i.e. 20 ns divided by 2.5 ns/clock). Ex. 2026 (Direct Rambus Tech. Disclosure), p. 9 ("In 100MHz SDRAMs systems, this change in Read/Write direction causes a 10ns clock cycle of delay. In contrast, a one cycle delay in Rambus is only 2.5ns")

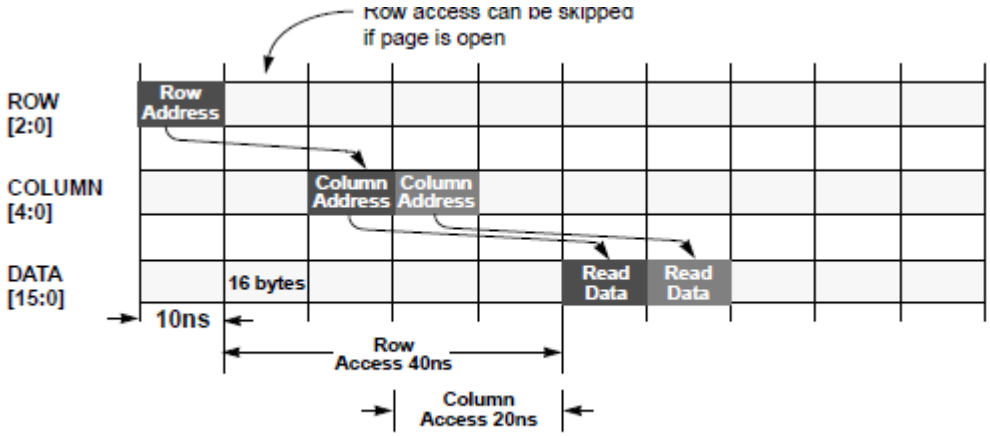

Figure 8: Read Transaction

Ex. 2026 (Direct Rambus Tech. Disclosure), p. 11.

182. But the JEDEC DDR2 SPD standard (Section 4.1.2.10 of JESD21-C) only enables encoding and communicating to the memory controller CAS latencies of 2 through 5.

Byte 18: SDRAM Device Attributes – $\overline{\text{CAS}}$ Latency

This byte describes which of the programmable $\overline{\text{CAS}}$ latencies are acceptable for the module. If the bit is "1", then that $\overline{\text{CAS}}$ latency is supported on the module; if the bit is "0", then that $\overline{\text{CAS}}$ latency is not supported by the module.

Bit 7	Bit 6	Bit 5	Bit 4	Bit 3	Bit 2	Bit 1	Bit 0
TBD	TBD	$\overline{\text{CAS}}$ Latency = 5	$\overline{\text{CAS}}$ Latency = 4	$\overline{\text{CAS}}$ Latency = 3	$\overline{\text{CAS}}$ Latency = 2	TBD	TBD
0	0	1 or 0	1 or 0	1 or 0	1 or 0	0	0
1 = Supported on this assembly; 0 = Not supported on this assembly.							

Ex. 2031 (JESD21-C, 4.1.2.10), p. 16.

183. Thus, the JEDEC SPD standard is incompatible with a module as contemplated by Perego, regardless of whether an extra cycle is added in the data path router 597, undermining Petition's statement that it would have been obvious to add a cycle to Perego so that the memory module complies with the timing of the JEDEC standards. What Petitioner is proposing as desirable would in fact be impossible. Petition, p. 95. The same is potentially true as well for DDR2 SDRAM embodiments of Perego. If the CL field of the SDRAMs is set to 5 as allowed by the DDR SDRAM standard, there is no SPD encoding to communicate a module latency of 6, as would be the case if there was additional latency due to the datapath router. There is no discussion of this issue within Perego. Ex. 1071, generally.

184. The Petition concludes this section by saying that a cycle would be added "so the "*circuitry*" has enough time to perform its functions (including "*register[ing]*" the data signals for interfaces 520a/b with latches 597f-m, blue, below) using "internal" clock circuit 570a-b (purple, below)."

185. The Petition recites 3 passages for Perego in support of these assertions: 18:65-19:3, 17:61-63, 12:65-13:5. None of these actually speaks to the question at hand. The passage at 12:65-13:5 is about the global clock source for the buffer device and not about the timing of data through the buffer device. The units 570a and 570b are clock circuits for generating and adjusting clock

signals so they meet phase requirements for proper operation of the buffer device. Ex. 1071, 12:12:52-13:5. They may include a "Direct Rambus Clock Generator" which was a device that Rambus specified and a variety of semiconductor vendors produced. Perego discloses that this functionality can be within the buffer device 405 instead of being a discrete separate integrated circuit. Regardless of its location, its functionality is the same. Ex. 2035 (Direct Rambus Clock Generator). Specifically it "includes signals to synchronize the Direct Rambus channel clock to an external system or processor clock." Ex. 2035 (Direct Rambus Clock Generator), p. 1. The nature of the Direct Rambus Channel is such that the phase of the ClockToMaster must be adjusted so that the clock signal arrives at the buffer device in phase alignment with the buffer-devices internal clock so that the incoming data can be properly captured. This is what the clock generator does. Ex. 2035 (Direct Rambus Clock Generator), p. 1 ("Synchronizes the Clock Domains of the Rambus Channel With An External System or Processor Clock."). The signal ClockToMaster and the buffer device internal clock are provided to the PCLKM and SyncClkN input puts to the clock generator and it adjusts the phase of CLK and CLKB so that those two inputs are phase aligned. This is all it does and it is only relevant to the Direct RDRAM-based embodiments. It is irrelevant to the question of whether there is an extra clock cycle of latency added to the path through the buffer circuit.

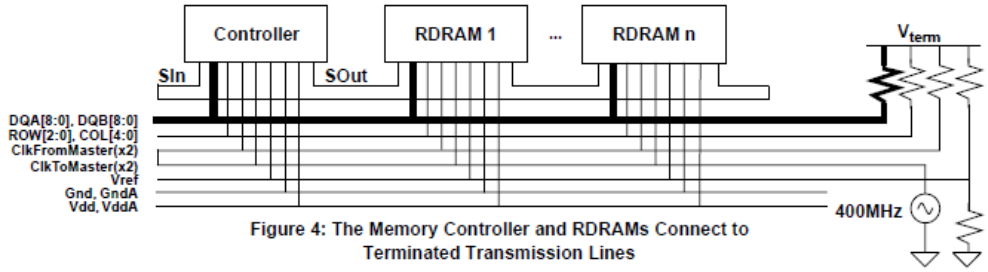

Figure 4: The Memory Controller and RDRAMs Connect to Terminated Transmission Lines

Ex. 2026 (Direct Rambus Tech. Disclosure), p. 8.

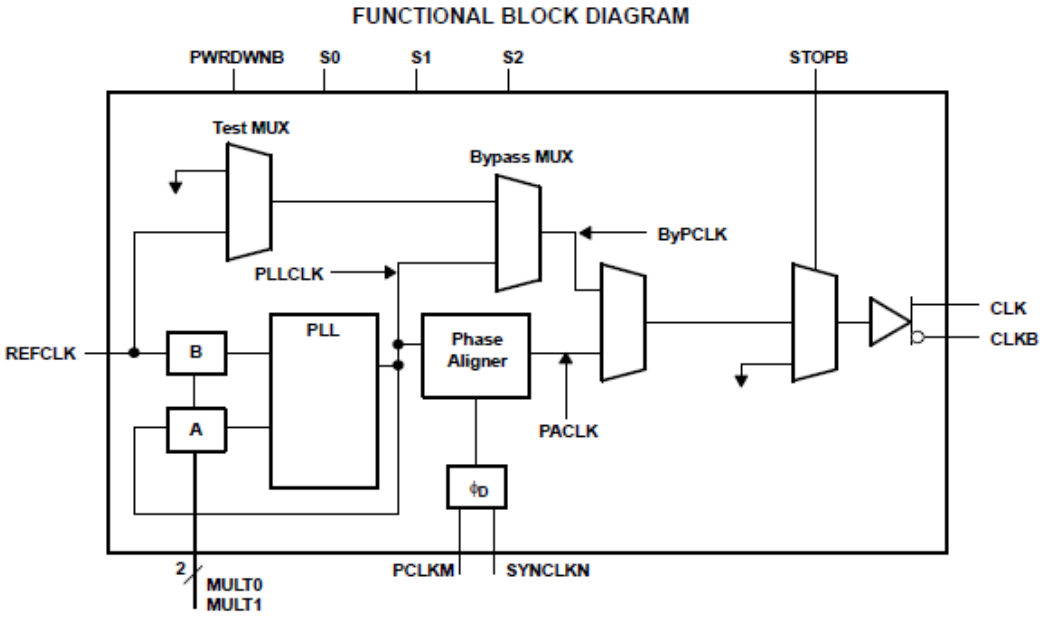

Ex. 2035 (Direct Rambus Clock Generator), p. 2.

186. The passage at 17:61-63 is about the structure of the multiplexer/demultiplexer, not its timing. And finally, the passage at 18:65-19:3 is about whether the multiplexer/demultiplexer 597 implements a full or partial crossbar. None of the support offered actually supports the statements made about the delay experienced by data as it traverses the buffer device.

187. I understand that Petitioner also contends that a prior Board finding regarding a different Netlist patent reciting a different limitation regarding CAS latency is binding on Netlist in this proceeding. Petition, pp. 93-94. I observe that the limitation at issue in that proceeding differs significantly from those recited in the '417 Patent, and a POSITA would not view them as co-extensive. See Ex. 1030 (Final Written Decision in IPR2017-00549), p. 3 (quoting claim 1 of U.S. Patent No. 8,756,364, which recites, in relevant part, "wherein the first logic element is configured to determine a latency value based on a previous memory command received by the memory module from the memory controller, and wherein the first logic element controls the second logic element to selectively enable the data communication according to the latency value"). For example, that limitation does not recite "wherein data transfers through the circuitry are registered for an amount of time delay such that the overall CAS latency of the memory module is greater than an actual operational CAS latency of each of the memory devices." '417 Patent, [1.f].

8. Summary

188. Claim 1 of the '417 Patent is a very detailed claim with many specific elements. They recite all aspects of a module organization and operation in a memory system. Petitioner has sought to invalidate this claim by principally pointing to the embodiments of Perego and by claiming that many of the

essential differences between these two references are inconsequential and/or obvious to overcome.

189. Specifically, the ’417 Patent discloses and claims modules that use JEDEC-compliant DRAMs and that have a memory bus interface that is similar to JEDEC-standardized RDIMMs. Improvements are made to improve the signal loading and improve bandwidth performance within the context of a multi-drop memory bus with conventional signal trace topologies.

190. In contrast, the principal and illustrated embodiments use unconventional DRAMs developed by Rambus Inc. and are not JEDEC-standardized. Conventional SDRAMs can also be used because of the protocol translation embodiments of the buffer device. But the essential and radically different approach of Perego is the focus on a dynamic point-to-point memory system topology and configurable data path widths allowing for changes to the capacity of the memory system, changing the number of used DQ signal traces on the primary channel without losing access to any of the storage locations.

191. And yet Petitioner, by misrepresenting the teaching of Perego, finds apparently obvious transformations to warp Perego’s approach into the very different disclosed and claimed modules of the ’417 Patent.

C. Dependent Claims 2 Through 15

192. All of the dependent claims of the '417 Patent depend on Claim 1. As each of the dependent claims directly or indirectly depend on Claim 1, Petitioner's failure to show the unpatentability of Claim 1 renders all of the Petition's arguments relating to Claims 2 through 15 moot.

1. Claim 5

193. Claim 5 recites "The memory module of claim 1, wherein the memory devices are organized in four ranks and the set of input address and control signals include four chip select signals, one for each of the four ranks". Petition, p. 101. The Petition attempts to find this element in Perego by looking to the illustrated embodiment of Figure 4B of Perego. But as explained above, this embodiment shows Rambus Direct RDRAMs. Rambus Direct RDRAMs do not receive a chip select signal. Therefore, according the Petitioner's own proposed construction, the colored groups of memory devices are not ranks. Further, the Petition says that the set of input address and control inputs as depicted in Figure 4B include four chip select signals. There is no support for this statement. There are no input chip select signals shown in Figure 4B and in general, and as shown above, the disclosed point-to-point primary channel of Perego does not include any chip select signals There are not four chip select signals, one for each of the memory devices on secondary channels 415a-d of

the embodiment of Figure 4B. Thus, regardless of the faulty analysis of claim 1, Petitioner has failed to show any of the additional portions of claim 5 to be present in Perego.

194. Thus, Petitioner has not shown that any of the challenged claims are disclosed or rendered obvious by Perego in combination with JESD79-2.

IX. GROUND 2 DOES NOT RENDER THE CHALLENGED CLAIMS OBVIOUS

195. I understand that, to the extent that Ellsberry is properly considered prior art to the '417 Patent (which I disagree with, as noted above), Ground 2 further relies on Ellsberry. However, this addition does not cure the deficiencies that I noted above for Ground 1, as I explain further below.

196. The Petition's analysis of Ground 2 begins with the analysis of Ground 1, which is based on the combination of Perego and JESD79-2 (the DDR2 SDRAM standard), and selectively adds aspects of Ellsberry. In introducing Ground 2, the Petition characterizes Ellsberry as being directed to "to efficiently organizing and using memory systems, including expanding memory module capacity without the need for additional chip select lines on the main memory bus." Petition, p. 113. It does this by teaching "rank multiplication." The Petition claims that this teaching of rank multiplication is a point of commonality between Perego and Ellsberry and therefore part of the motivation to combine Ellsberry with the references of Ground 1. But as

explained above, the concept of rank multiplication is completely foreign to Perego, and so this is not a point or commonality but rather a teaching that differentiates and separates the two references.

197. The Petition further justifies the combination of Ground 1 with Ellsberry based on a number of additional flawed arguments outlining alleged similarities between Perego and Ellsberry.

198. For example, the Petitioner claim that the two references disclose modules of similar structure. Petition pp. 111-113. However, the coloring of Figure 3C of Perego suggests similarities that do not exist. As described supra, the module of Figure 3C is agnostic as to what type of synchronous DRAMs are being used. They can be DDR2 SDRAMs, Direct RDRAMs, or XDR RDRAMs, or some other variety of memory type. Ex. 1071, 8:1-5. A POSITA would understand that the placement of the configurable width buffer device 391 in the middle of the memory module is a structure that minimizes trace lengths and improves signal integrity and is not indicative the two groups of memory devices 360 being logically grouped into two ranks or other organizations. Indeed, with respect to the disclosed embodiments that illustrate this, the internals of the configurable width buffer device 391 (i.e., Figure 5B, below), the two interface units 520a and 520b are together and identically controlled by the singular request and address logic unit 540. In the embodiments disclosed that could use

JEDEC-Standardized SDRAMs, a POSITA would understand that all the SDRAMs are collectively part of a single rank.

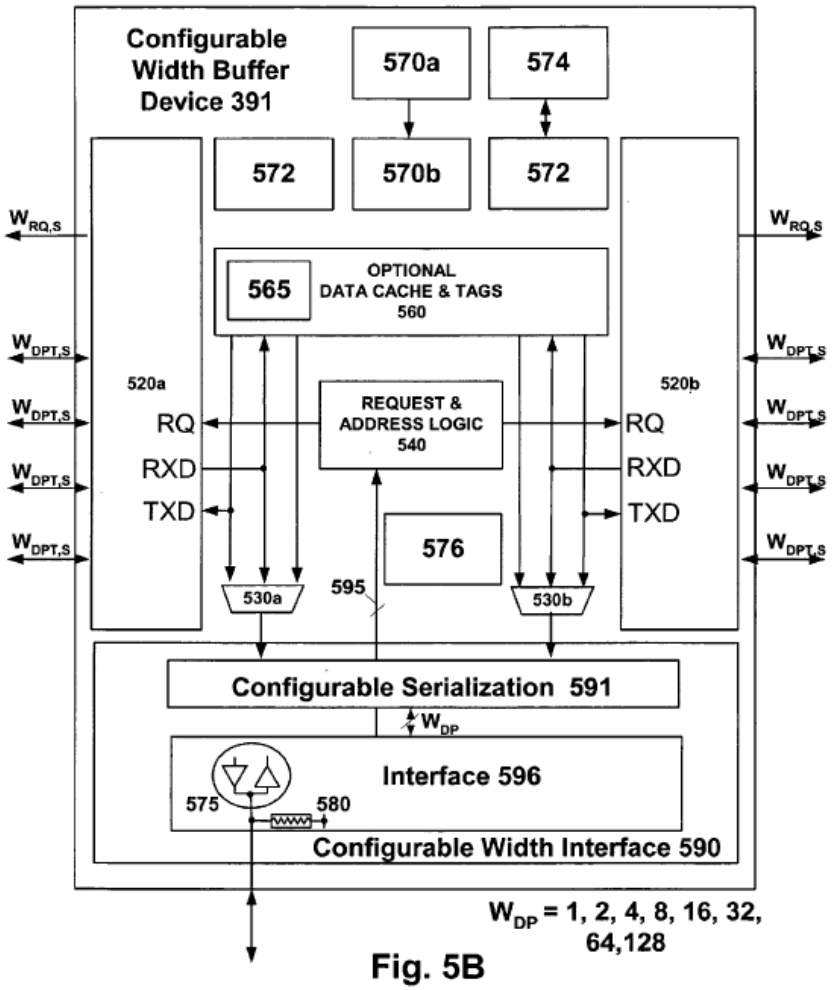

Fig. 5B

Ex. 1071, Figure 5B.

199. In contrast, the embodiment of Figure 10 of Ellsberry "illustrates a single chip-select memory configuration in which one control unit 1002 and one bank switch 1004 are used to control two memory banks 1006 & 1008, each memory bank having two memory devices 1010 & 1012 in separate data buses

1014 & 1016." Ex. 1073, [0053]. This embodiment discloses an x8 module of fixed width and singular memory type and with two ranks. It has very little in common with the embodiment of Figure 3C of Perego.

200. Petitioner also claims that the relevance of the JEDEC-standards, including the DDR2 SDRAM standard JESD79-2 to both Perego and Ellsberry is a commonality motivating the combination. Petition, p. 114. This also a distortion of the teachings of both references. The memory devices defined by JESD79 and JESD79-2 are the only memory devices used by Ellsberry in an embodiment, and the module interface of its principal embodiments are closely aligned with the JEDEC-standardized modules such as those of Section 4.10.20 of JESD21-C. Ex. 1073, [0008], [0009], [0026], [0046].

201. In contrast, the modules of Perego's disclosed embodiments are profoundly different. They principally use DRAMs defined by Rambus and not JEDEC, and Perego's embodiment, such as that of Figure 2A, that could use JEDEC-standardized memories on the module use protocol translation in the buffer device to present an interface that is very different and bear very little resemblance to JEDEC-defined DIMM modules. Ex. 1071, Figures 2A, 3C, 4A, 4B, 4C, 5B, 2:8-10, 9:53-57, 10:63-67.

202. Finally, the Petition suggests commonality exists between the references because they both buffer signals on the module, advantageously

reducing the load on the data lines. Petition, pp. 114-115. This too is misleading. Perego's approach to signal integrity on the primary channel is principally focused on changing the memory systems interconnect architecture. A POSITA would recognize that Perego's memory system organization of Figure 3A and 3B and the use of a Rambus Channel as the principal interconnect as in the embodiment of Figure 2A are representative of Perego's approach and that this approach is much more comprehensive change to the memory system architecture than the simple buffering of the DQ signals on a module. Ex. 1073, Figures 5, 6, [0012], [0047].

203. But the essential problem with the articulated combination of Ellsberry's rank multiplication into Perego as described by the Petition and Dr. Wolfe is that it represents a big step backwards, one that a POSITA would not be motivated to take. Specifically, Dr. Wolfe describes the result of the combination as a module with the bank switch of Ellsberry instead of the point-to-point interface 510 or the configurable width interface 590 in order to yield a "memory module [that] is dual inline memory module compl[ia]nt with a Joint Electron Device Engineering Counsel (JEDEC) standard." Ex. 1003, ¶ 196. Without these either of these high performance interfaces and the bandwidth concentrating multiplexers 530a,b, much of the advantage and uniqueness of the Perego memory system architecture would be lost. Ex. 1071, Figure 5A, 5B,

Among the critical benefits of Perego that would be lost would be the configurable primary channel width, the point-to-point primary bus interface, the ability to have "different types of memory devices on different modules within a memory system by employing buffer device to translate protocols employed by controller 310 to the protocol utilized in a particular memory device implementation." Ex. 1071, Figure 5A, Figure 5B, 9:50-57, 10:54-67, 11:1-25, 11:56-12:3. This benefit is key because it allows the use of memory devices that are cheapest or provide the best cost/performance applicable to a particular system, yielding either performance levels or cost savings that are not available to a memory system compliant with a JEDEC standard. These cost savings and tradeoffs would be available to a system architect at all times, not just the short-lived benefits that rank multiplication delivers at the beginning of the life of a new generation of higher density memory devices for the brief period when the new generation is available but at a higher cost per bit than that of the previous (lower capacity) generation.

204. Because I do not agree that a POSITA would have been motivated to combine Ellsberry with Perego and JESD79-2, I also disagree that Ground 2 renders any of the challenged claims obvious.

205. The Petition's analysis of Ground 2 begins by incorporating the analysis of Ground 1 without any input or modification from Ellsberry. Petition,

p. 115. But as shown above, the analysis provided for Ground 1 is based on a pervasive misunderstanding of the teaching of Perego and the details of the relied-upon embodiments. These flaws are therefore brought into the analysis of Ground 2 and the Ground 2 analysis of claims 1-15 fail for the same reasons as exposed above.

206. The Petition then addresses Claims 1-3, 6, 8, and 10-11, looking to the rank multiplication of Ellsberry to augment the teaching of Perego and JESD79-2. But Petitioner believes that Perego already teaches rank multiplication, so it is unclear how two teachings of rank multiplication are better or more effective than one at disclosing the claims of the '417 Patent, which are not directed to rank multiplication as discussed above.

207. Specifically, Petitioner applies the teaching of Ellseberry to only elements [1.e.3] (relating to the overall CAS latency of the memory module), [1.f] (wherein data transfers are registered for an amount of time), [9] (the memory module comprising a phase locked look clock driver), [4] (x64 or x72 ranks constructed of x4 memories), [13] (load isolation), [14] (a specific data rate for the module and for the transfer through the data paths), and [15] (relating specifically to write operations). Petition, pp. 118-120. The analyses for all the other elements in all the other claims are those from Ground 1 without modification. But as explained at length above, the analysis of Ground 1 for

Claim 1 and consequently all the dependent claims is flawed and does not render the claims of the '417 Patent obvious. The same flaws and conclusion thus also apply to Ground 2. Most notably, Petitioner has failed to show disclosure of the chip select and rank elements of Claim 1 by the combination of references offered in Ground 1 and consequently the same result holds for Ground 2 as well.

X. GROUND 3 DOES NOT RENDER THE CHALLENGED CLAIMS OBVIOUS

208. I understand that Ground 3 further relies on Halbert, in combination with Perego and JESD79-2. However, this addition does not cure the deficiencies that I noted above for Ground 1, as I explain further below.

209. The Petition's analysis of Ground 3 begins with the analysis of Ground 1, which is based on the combination of Perego and JESD79-2 (the DDR2 SDRAM standard), and selectively adds aspects of Halbert. In introducing Ground 3, the Petition characterizes Halbert as being directed "to efficiently organizing and using memory systems, including expanding DIMM-formatted memory modules by using one or more buffers to isolate the memory devices from the memory controller, using DDR SDRAM memory devices." Petition, pp. 121-122. The Petition characterizes DDR SDRAM memory devices, which are defined by the JEDEC standard JESD79, not JESD79-2, as being similar to the DDR2 SDRAM memory devices relied upon in Ground 1. The Petition fails to identify the similarities and/or differences between DDR

SDRAMs and DDR2 SDRAMs and how the differences might affect their analysis.

210. The Petition's analysis of Ground 3 begins by incorporating the analysis of Ground 1 without any input or modification from Halbert. Petition, p. 123. But as shown above\, the analysis provided for Ground 1 is based on a pervasive misunderstanding of the teaching of Perego and the details of the relied-upon embodiments. These flaws are therefore brought into the analysis of Ground 3 and the Ground 3 analysis of claims 1-15 fail for the same reasons as exposed above.

211. Petitioner then goes on to apply Halbert to only elements [1.d.1] (N-bit wide ranks), [1.e.3] (relating to the overall CAS latency of the memory module), [1.f] (wherein data transfers are registered for an amount of time), [9] (the memory module comprising a phase locked look clock driver), and [2] (load isolation). Petition, pp. 123-126. The analyses for all the other elements in all the other claims are those from Ground 1 without modification. But as explained at length above, the analysis of Ground 1 for Claim 1 and consequently all the dependent claims is flawed and does not render the claims of the '417 Patent unpatentable. The same flaws and conclusion thus also apply to Ground 3.

212. The Petition's understanding of Halbert is flawed to the same extent as it's analysis of Perego. Most notably, the two "ranks" 140 and 142 of Figure 4

are together part of a singular memory device array that receives a single registered address and command bus (RADD/RCMD). These two subunits 140 and 142 are always shown to behave identically to the receipt of commands on this bus. Ex. 1078, Figure 4, Figure 5, Figure 6. This singular RADD/RCMD bus would have one CS signal as part of the singular registered command bus. To the extent that is it argued, regardless of what is shown in Figure 4, that the two units 140 and 142 receive unique commands on separate signals, because the two collections of memory devices would always receive CS signals that are always the same value because they are shown in Figures 5 and 6 as always behaving identically. They would always at the same time both be active or both be inactive. Therefore Halbert does not disclose any claim element that requires a plurality of registered chip select signals including one registered chip select signal having an active signal value and one or more other registered chip select signals each having a non-active signal value, as, for example, Claim 1 and all of the dependent claims of the ’417 Patent require.

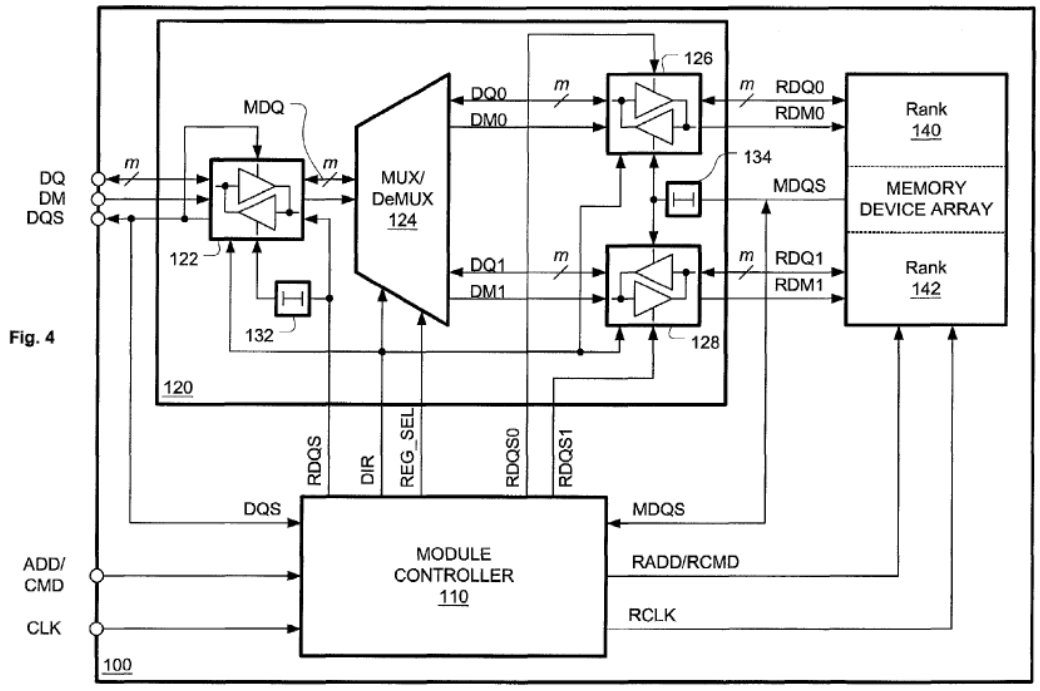

Ex. 1078, Figure 4.

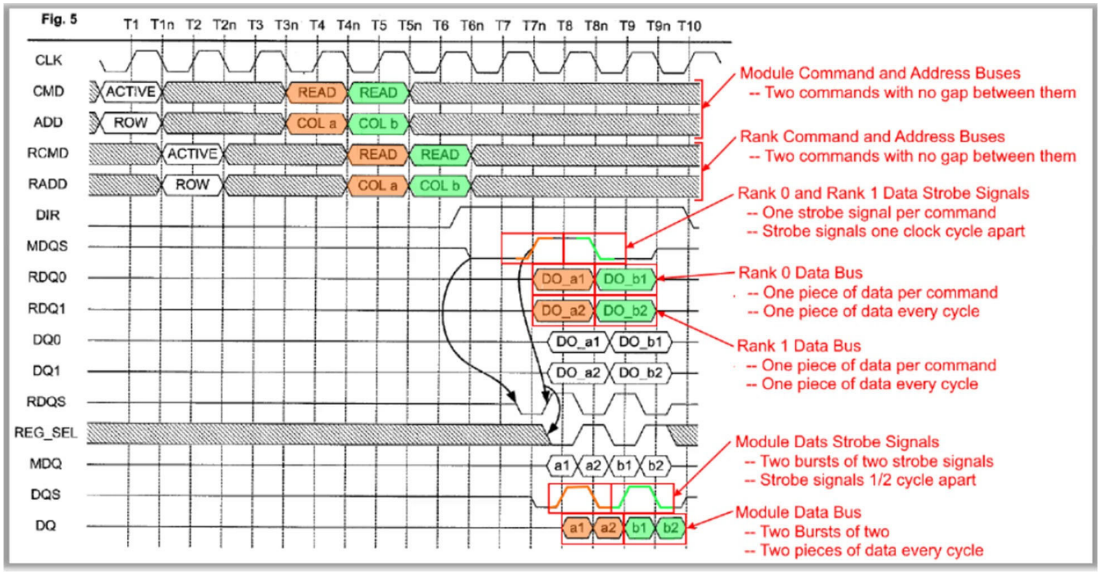

Ex. 1078, Figure 5, Annotated.

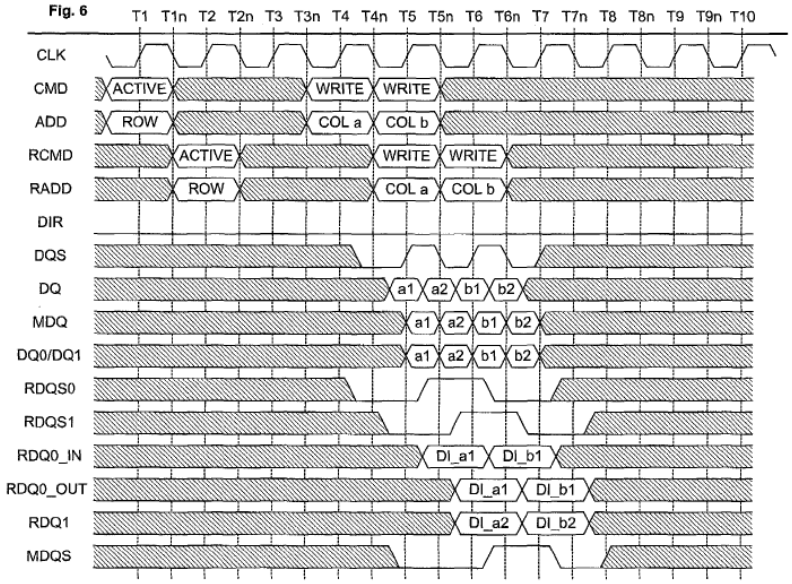

Ex. 1078, Figure 6.

213. There is no support for the idea that memory devices in 140 and those in 142 receive separate or different registered chip select signals. They are in fact the two halves of a single 2m-wide memory structure. The two halves of this single structure are interleaved within data interface circuit 120. The function of this data interface circuit 120 is to transform a data stream that is 2m bits wide with one piece of data every cycle into and from a data stream that is m bits wide with two pieces of data every cycle. Specifically, Halbert aims to "widen[] the data bus on the memory module as compared to the width of the system memory data bus," as this "allow[s] a faster system memory data bus to operate at full speed with slower memory devices." Ex. 1078, 3:43-48. According to Petitioner's proposed construction, all the memory devices in the memory device array constitute a single rank.

214. To the extent that Petitioner tries to change its argument by pointing to the two select signals in the prior art Figure 2, it is important to note that this embodiment is distinct from and not combinable with the embodiment of Figure 4. It does not include a buffer circuit and data transactions involve only one of the two banks at a time: it is incompatible with the stated goal of Halbert to "allow[s] a faster system memory data bus to operate at full speed with slower memory devices." Ex. 1078, 3:43-48. Even Dr. Wolfe agreed that "ranks" 140 and 142 collectively constitute a single rank under Petitioner's proposed construction. Ex. 2033 (Wolfe Deposition Transcript), 162:4-163:23. He later appears to have attempted to hedge that statement after being shown Halbert's statement that "[g]enerally, multiple ranks will receive the same address and commands, and will perform memory operations with the interface circuit concurrently," Ex. 1078, 4:57-59, but he testified that they would be two ranks "if they receive different commands." Ex. 2033 (Wolfe Deposition Transcript), 179:20-180:16. The Petition, however, identifies no such disclosure in Halbert wherein "ranks" 401 and 402 receive different commands. See Petition 120-126.

215. Although Halbert does say that DDR SDRAMs can be used in the present system (Ex. 1078, 9:55-59), the memory devices illustrated in Figures 4, 5, and 6 and described in the corresponding text are distinctly not JEDEC-compliant DDR SDRAM. Specifically, like DDR2 SDRAMs, DDR SDRAMs

transfer two pieces of data in the time that a single command is communicated. In Figure 5, for example, the period of a command on the RADD/RCMD bus is one clock cycle. In this time, the DRAMs only transmitted one piece of data, not two. Petitioner's glossing over these issues by noting the similarity between DDR SDRAMs and DDR2 SDRAMs is irrelevant to understanding the teaching and applicability of Halbert in the present matter.

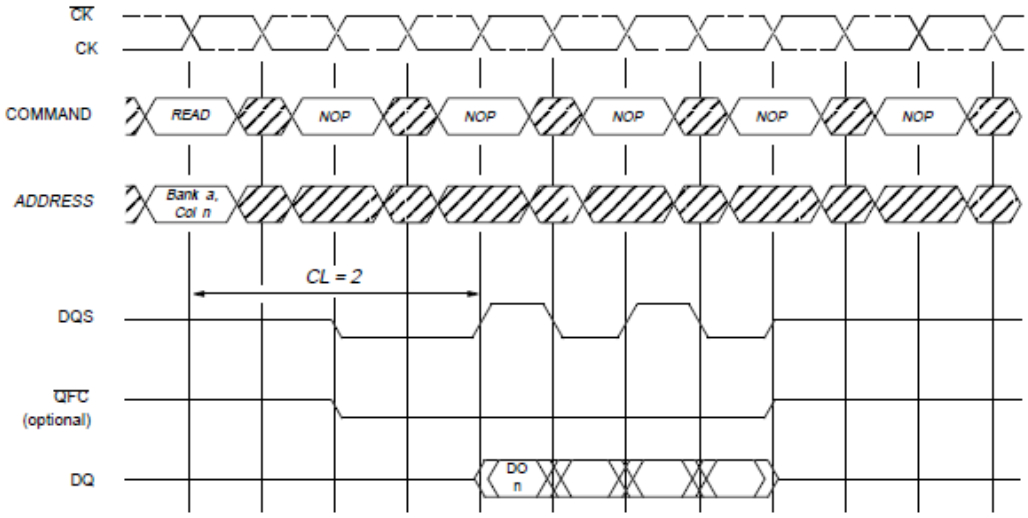

Ex. 1069 (JESD79, defining DDR SDRAMs), Figure 7.

216. The embodiments of Halbert referred to in the Petition do not support the teaching of a plurality of N-bit ranks and or a plurality of input chip select signals or registered chip select signals. Thus, Halbert cannot cure these and other deficiencies in Perego identified with respect to Ground 1. Therefore Ground 3 also fails to render the claims of the '417 Patent obvious.

217. Petitioner particularly fails to identify where Halbert discloses limitations [1.e.3], [1.f], and [9] related to data transfer delays and CAS latency. Petitioner merely asserts that Figures 3 and 5 of Halbert render obvious limitations "concerning '*overall CAS latency of the memory module*' that is '*greater than an actual operational CAS latency of each of the memory devices*' by '*at least one clock cycle.*'" Petition, pp. 124-126. The annotated Figure 3 of Halbert on page 125 of the Petition makes clear that the Petitioner is relying on the delay between T4 and T5—i.e., the delay in the control path between a command being presented on the ADD/CMD bus and the corresponding command emerging from Halbert's Module Controller 110 onto RADD/RCMD as the source of the increased module CAS latency. Element [1.f], however, recites in full "wherein data transfers through the circuitry are registered for an amount of time delay such that the overall CAS latency of the memory module is greater than an actual operational CAS latency of each of the memory devices." As explained above in connection with Perego and JESD79-2, this element explicitly requires data transfers to be registered through the claimed circuitry en route between the memory devices and the primary link. Thus, the claimed time delay must include delay in the data path. Delay on the address/control path alone cannot satisfy this element. Therefore, as with JESCD21-C, Dr. Wolfe's reliance on Halbert "registering the address and

command signals" and the resultant delay (e.g., Ex. 1003, ¶ 366 (citing Ex. 1078, 2:46-55)) is irrelevant to this claim element.

218.

XI. NO SIMULTANEOUS INVENTION

219. I understand that Petitioner asserts that Ellsberry, Perego, an additional reference not asserted in this IPR, as well as other patents related to the '417 Patent, demonstrate "simultaneous invention" of "rank multiplication." Petition, pp. 126-127. I understand that simultaneous invention requires the actual invention to cover the complete claims. The fact that none of the references are even alleged to be anticipatory means there is no simultaneous invention.

220. Further, Petitioner asserts that these references and patents "recognized and solved the same problem" as the '417 Patent through "rank multiplication." *Id.* As explained above, the '417 Patent does ***not*** claim rank multiplication, and Perego does not even disclose rank multiplication. Moreover, the Petition does not establish any contemporaneous invention of the subject matter actually claimed in the '417 Patent.

* * *

I declare that all statements made herein of my knowledge are true, and that all statements made on information and belief are believed to be true, and that all these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code.

Executed this 8th Day of December, 2023.

Steven Przybylski, Ph.D.

APPENDIX A

Steven A. Przybylski

President, Verdande Group, Inc.
5630 Meadow Lane, Ann Arbor MI 48105
V: 734-484-3574 F: 734-484-3390
sp@verdande.com

Professional Experience

2003 - Present **President,** Verdande Group, Inc

An inventor, registered patent agent, strong communicator and seasoned technical expert, I provide intellectual property services to a diverse set of clients across the digital electronics spectrum. I provide patent, trade secret, and contract dispute litigation support services, and have extensive experience as a testifying and non-testifying expert in matters before district courts, the PTAB, the ITC, and foreign jurisdictions. I have testified by expert report, deposition, and live testimony on matters of claim construction, infringement, and invalidity, and have assisted in evaluations of potential damages and monetization strategies. With strong written and verbal skills, I provide tutorials to those of both technical and non-technical backgrounds, including judges and attorneys. In addition, I offer pre- and post-transaction portfolio evaluations and due diligence analyses.

I have been retained for these services with regard to:

Computer Systems	Consumer Electronics	Semiconductors
Processor design	Infotainment systems	DRAMs
Processor architecture	Video encoding and decoding	Flash
Multiprocessing	Medical equipment	SRAMs
Graphics systems	Embedded systems	Other non-volatile memories
Graphics/CUDA processors	Portable electronic systems	Memory controllers
Cache memories	User interface design	High-speed interconnects
Memory systems	Wired and wireless networking	Signal integrity
Memory modules	Power/energy management	2D, 2.5D and 3D packaging
Floating point computation	Standards development	Error detection and correction
Operating systems	Display interconnects	Clock distribution techniques
I/O systems	Integrated memories	Semiconductor processing
System boot issues	Digital signal processing	IC and system testing
	Cloud-based content navigation and serving	Computer aided design
		FPGAs

2000 – 2003 **Manager,** SSM Technologies, LLC.

I managed, developed and sold a significant patent portfolio relating to DRAMs. Recognizing the intrinsic value of a long-neglected portfolio, I negotiated control of the portfolio and developed a new patent prosecution strategy. I oversaw and participated in the implementation of this strategy, culminating in the issue of a new set of broad method and apparatus claims. I then worked to monetarize the portfolio and to maximize return to the inventors and SSM's investors by analyzing the sell/license decision and guiding the stakeholders through the decision-making process. I gathered relevant documentation and analyses, approached likely purchasers, presented interested parties with the collected materials, actively sold the offensive and defensive utility of the portfolio, and orchestrated an auction-style final sales process.

1991 – 2000 **President and Principal Consultant,** Verdande Group, Inc

I have consulted on product strategy, comparative technical analysis, and system-level design in the areas of DRAMs, Flash, SRAMs and other semiconductor memories, computer systems and CPU architecture, and uni- and multi-processor memory hierarchy design. Technical engagements have included serving as memory system architect for Flash and DRAM memory devices and systems. In addition, I offered: general business planning and analysis services to computer and semiconductor companies; intellectual property management, including portfolio evaluation, development and brokering, and litigation support services to IP holders and legal firms; and due-diligence analyses to venture capitalists and securities firms.

My writing and analysis encompasses an extensive range of topics for trade magazines, conferences, and academic journals. As one of the leading independent analysts of the semiconductor memory industry, I wrote (1994) and substantially revised (1996) an 850-page research report entitled *New DRAM Technologies: A Comprehensive Analysis of the New Architectures*, available from the publishers of *The Microprocessor Report*. I have also taught and presented tutorials to both technical and non-technical audiences on subjects including semiconductor memory architecture, memory hierarchy design, and computer systems and CPU architecture. In particular, I have presented over 50 full- and partial-day seminars on these topics to over 2,500 managers, engineers, marketers, and analysts. I have also served on the advisory boards of high-technology start-ups, advising on processor, DRAM and Flash memory trends and general business opportunities.

1989 - 1991 **Chief Scientist and Systems Architect,** MIPS Computer Systems.
I played a variety of roles within MIPS, including assistant to the Senior Vice President of Engineering, technical lead for a group of designers that investigated the technical and economic feasibility of developing a companion chip to the R3000 CPU and R3010 FPU, troubleshooter within the High-End Systems Group responsible for taking on any short or long term task that was in need of additional manpower, chief scientist for the group in charge of planning and investigating options for future high-end systems, systems architect for both a new R4000 based symmetrical multi-processor and a low-cost desk-top machine.

1990 - 1991 **Consulting Assistant Professor,** Dept. of Elect. Eng., Stanford University.
I designed and taught an advanced graduate level course on cache and memory hierarchy design and supervised the students' research activities. Two groups of students continued their research and ultimately published their efforts under my direction.

1989 - 1989 **Postdoctoral Scholar,** Department of Electrical Engineering, Stanford University.
As a postdoctoral scholar, I continued my dissertation research into the analysis of cache behaviour and performance-optimal memory hierarchy design.

1984 - 1985 **Computer Architect and Systems Designer,** MIPS Computer Systems.
My main contribution during MIPS' first year was the specification of the architecture and micro-organization of the R2000 VLSI RISC CPU, including the CPU interface and the cache sub-system. I was integrally involved in the implementation of both the CPU and system-level design. My most important role was that of liaison between the VLSI, board-level hardware, compiler and operating system teams.

1981 - 1988 **Research Assistant,** Department of Electrical Engineering, Stanford University.
I was a key participant in the MIPS VLSI CPU design team, involved in the design, debugging and testing of the IC design. I also oversaw the development of a prototype CPU board, and subsequently aided in the architectural specification of the MIPS-X VLSI CPU. My dissertation research used analytical and empirical techniques to examine memory hierarchy design from the perspective of optimizing system level performance.

Education

Haas School of Business, University of California at Berkeley, M.B.A. (2000).
Coursework and final project included a combined emphasis on business strategy and financial analysis.

Stanford University, M.S.E.E. (1982), Ph.D. (1988).
Research areas: Computer systems design and VLSI integrated circuit design. Concentrated areas of study spanned the breadth of both electrical engineering and software systems. Topics of extensive coursework included digital systems design, computer architecture, semiconductor devices, error correction codes, control theory, and system software — including user interface design, operating systems, compilers, and I/O systems. Doctoral dissertation title: Performance-Directed Memory Hierarchy Design.

University of Toronto, B.A.Sc. with Honours (1980).
My unique, specially approved course of study within the Engineering Science Programme combined the computer science and electrical engineering degree programs. This provided a broad and deep background in both areas. Coursework and undergraduate thesis covered digital systems design, computer systems, control theory, network analysis, analog circuit design, semiconductor physics, human-computer interfaces, operating systems, compilers, database design, and AI.

Honors and Publications

Awards

Most Influential Paper Award, ACM/IEEE International Symposium on Computer Architecture, 1989, with Mark Horowitz and John Hennessy.

Micro Test-of-Time Award, for one of the ten most influential papers of the first 25 years of the IEEE Micro Symposium, with John Hennessy, Norman Jouppi, Christopher Rowen, Thomas Gross, Forest Baskett, and John Gill. 1994.

Books

S. Przybylski, *New DRAM Technologies*, MicroDesign Resources, 1996, 1994.

S. Przybylski. *Cache and Memory Hierarchy Design: A Performance-Directed Approach*, Morgan Kaufmann, 1990.

J. Hennessy, S. Przybylski. VLSI Processor Design Methodology. Chapter I in *VLSI Electronics, Microstructure Science*, Vol. 14, Einspruch, N. ed., Academic Press, 1986.

Journal Articles

T. Wada, et al. An Analytical Access Time Model For On-Chip Cache Memories, *Journal of Solid-State Circuits*, Vol. 27, No. 8, August 1992.

T. Gross, et al. Measurement and Evaluation of the MIPS Architecture and Processor. *ACM Transactions on Computer Systems*, Vol. 6, No. 3, August 1988.

M. Horowitz, et al. MIPS-X: A 20 MIPS Peak, 32-bit Microprocessor with On-Chip Cache. *Journal of Solid-State Circuits*, Vol. 22, No. 5, October 1987.

C. Rowen, et al. RISC VLSI Design for System-Level Performance. *VLSI Systems Design*, Vol. 7, No. 3, March 1986.

S. Przybylski, et al. The Organization and VLSI Implementation of MIPS. *Journal of VLSI and Computer Systems*, Vol. 1, No. 2, December 1984.

Other Articles

T. Gross, N. Jouppi, J. Hennessy, S. Przybylski, C. Rowen. A Retrospective on "MIPS: A Microprocessor Architecture", *IEEE Micro*, Vol. 36, No. 4, July/August, 2016.

R. Schuetz, H.J. Oh, J.-K. Kim, H.-B. Pyeon, S. Przybylski, P. Gillingham. HyperLink NAND Flash Architecture for Mass Storage Applications, 22nd IEEE Nonvolatile Semiconductor Memory Workshop (NVSMW), August 2007.

S. Przybylski, Intel's RDRAM Strategy a Sure Winner, *Microprocessor Report*, Vol. 11, No. 5, April 21, 1997

S. Przybylski, Intel's Gambles on a Sure Thing, *Electronic Engineering Times*, No. 947, March 31, 1997.

S. Przybylski, Will 1997 be the year of the 64-Mbit DRAM?, *Computer Design*, March 1997.

S. Przybylski, Combining DRAM and System Architecture, *OEM Magazine*, Vol. 4, No. 28, June, 1996.

S. Przybylski, SDRAMs Ready to Enter Mainstream, *Microprocessor Report*, Vol. 10, No. 6, May 6, 1996.

S. Przybylski, MoSys Reveals MDRAM Architecture, *Microprocessor Report*, Vol. 9, No. 17, December 25, 1995.

S. Przybylski, New DRAMs Improve Bandwidth, *Microprocessor Report*, Vol. 7, No. 2, February 15, 1993.

S. Przybylski, DRAMs for New Memory Systems *Microprocessor Report*, Vol. 7, No. 3, March 8, 1993.

S. Przybylski, DRAMs for New Memory Systems, *Microprocessor Report*, Vol. 7, No. 4, March 29, 1993.

S. Przybylski. The Performance Impact of Block Sizes and Fetch Strategies. 17th Annual International Symposium on Computer Architecture, May 1990.

S. Przybylski, J. Hennessy, M. Horowitz. Characteristics of Performance-Optimal Multi-Level Cache Hierarchies. 16th Annual International Symposium on Computer Architecture, June 1989.

S. Przybylski, M. Horowitz, J. Hennessy. Performance Tradeoffs in Cache Design. 15th Annual International Symposium on Computer Architecture. June 1988.

M. Horowitz, et al. A 12 MIPS 32b Microprocessor with On-Chip Cache. International Solid State Circuits Conference, February 1987.

J. Moussouris, et al. A RISC Microprocessor with Integral MMU and Cache Interface. International Conference on Computer Design (ICCD), October 1986.

J. Moussouris, et al. A CMOS RISC Processor with Integrated System Functions. Proceedings of COMPCON, Spring 1986.

C. Rowen, et al. A Pipelined 32b nMOS Microprocessor. International Solid State Circuits Conference, February 1984.

S. Przybylski. The Design Verification and Testing of MIPS. MIT Conference on Advanced Research in VLSI, January 1984.

J. Hennessy, et al. Performance Issues in VLSI Processor Design. International Conference on Computer Design (ICCD), November 1983.

J. Hennessy, et al. Design of a High Performance VLSI Processor. Third Caltech Conference on VLSI, March 1983.

J. Hennessy, et al. MIPS: A Microprocessor Architecture. MICRO-15, October 1982.

Technical Reports

S. Przybylski. *Performance-Directed Memory Hierarchy Design.* Technical Report No. CSL-TR-88-366, Stanford University, September 1988.

S. Przybylski. *The Implementation of MIPS.* Internal Technical Report, Computer Systems Laboratory, Stanford University, August 1984.

J. Gill, et al. *Summary of MIPS Instruction.* CSL Technical Note No. 237, Computer Systems Laboratory, Stanford University, March 1983.

J. Hennessy, et al. *Hardware/Software Tradeoffs for Increased Performance.* CSL Technical Report No. 228, Computer Systems Laboratory, Stanford University, February 1983.

Seminars, Tutorial, Session Chairs

PC System Architecture Trends, Panel member, 1999 Microprocessor Forum, San Jose, October 1999.

An Introduction to Semiconductor Memories, A Verdande Group presentation in cooperation with Citibank Inc, Taipei, Taiwan ROC, June 1998.

Embedded DRAMs: Today and Toward System-Level Integration, Half-day seminar. Presented more than 10 times to corporations between September 1997 and October 1999

The New DRAMs for Main Memory and Graphics, Full day Seminar, Verdande Group Seminar Series, Portland, Austin, San Jose, and Boston, September 1997.

Embedded DRAMs: Today and Toward System-Level Integration, Half-day seminar, Verdande Group Seminar Series, Portland, Austin, San Jose, and Boston, September 1997.

Sorting Out the New DRAMs, Half-Day Tutorial, Hot Chips IX, August 1997.

Embedded DRAMs: Today and Toward System-Level Integration, Half-day seminar, Presented at 24th Annual International Symposium on Computer Architecture, May 1997.

The New DRAMs For Main Memory and Graphics, Full day tutorial, 1997 PC Tech Forum, May 1997.

The New DRAM Architectures, DRAM Tutorial, 1996 ISSCC, February 1997.

Selecting the Best Memory Technology for PCs and Graphics, Full day tutorial, 1996 PC Tech Forum, May 1996.

The Changing DRAM and VRAM Marketplace, Full day tutorial, Presented more than 20 times to corporations between September 1995 and October 1999.

The Changing DRAM and VRAM Marketplace, Full day tutorial, Sponsored by Nikkei Microdevices, Tokyo, Japan, March 1996.

The Changing DRAM and VRAM Marketplace, Full day tutorial, 1995 Microprocessor Forum, October 1995.

Cache and Main Memory Options. Lecture. WinHEC, March 1995.

New Graphics Memories. Lecture. WinHEC, March 1995.

Emerging DRAM Architectures. Lecture. WinHEC, February 1994.

New DRAM Organizations. Session chair. Proceedings of COMPCON, Spring 1994.

Computer architecture and organization. Half-day tutorial. OCATE/Oregon State University, August 18, 1993.

DRAM Choices for the 1990s. Half-day tutorial. Hot Chips V, August 1993.

Computer Architecture. Full day tutorial with Normal Jouppi and Mark Horowitz. International Solid State Circuits Conference, February 1993.

Recent Trends in Processor Design, Reclimbing the Complexity Curve. Week-long course with Mark Horowitz and Mike Smith. Western Institute of Computer Science/Stanford University, August 1992 and August 1993.

Tracking the Processor-Memory Bandwidth Problem. Session chair, Microprocessor Forum, October 1992.

Patents

H. Pyeon, H. Oh, J.-K. Kim, S. Przybylski, *ID generation apparatus and method for serially interconnected devices*, Patent No. 8,984,249, March 17, 2015.

S. Przybylski, *Ring-of-clusters network topologies*, Patent No. 8,902,910, December 2, 2014.

S. Przybylski, R. Schuetz, H. Oh, H.B.Pyeon, *System having one or more memory devices*, Patent No. 8,812,768, August 19, 2014.

S. Przybylski, *Error detection and correction codes for channels and memories with incomplete error characteristics*, Patent 8,806,305, August 12, 2014.

J.-K. Kim, H. Oh, H.B. Pyeon, S. Przybylski. *Scalable Memory System*. Patent No. 8,671,252, March 11, 2014.

S. Przybylski, *Ring-of-clusters network topologies*, Patent No. 8,594,110, November 26, 2013.

S. Przybylski, *Error detection and correction codes for channels and memories with incomplete error characteristics*, Patent No. 8,429,495, April 23, 2013.

J.-K. Kim, H. Oh, H. Pyeon, S. Przybylski. *Scalable Memory System*. Patent No. 8,407,395, March 26, 2013.

L. Crudele, J. Moussouris, S. Przybylski. *System Having An Address Generating Unit and a Log Comparator Packaged as an Integrated Circuit Separate From Cache Log Memory and Cache Data Memory*. Patent No. 5,113,506, May 12, 1992.

L. Crudele, J. Moussouris, S. Przybylski. *CPU Chip Having Tag Comparator and Address Translation Unit on Chip and Connected to Off-Chip Cache and Main Memories*. Patent No. 4,953,073, August 28, 1990.

C. Hansen, S. Przybylski, T, Riordan, L. Weber. *Dual Byte Order Computer Architecture.* Patent No. 4,959,779, September 25, 1990.

Foreign counterparts of the above.

Additional Information

Registered US Patent Agent, 1997 - present

Professional Societies: IEEE, ACM

Citizenship: Canadian with U.S. resident alien status